

Scalable SUM-Shrinkage Schemes for Distributed Monitoring Large-Scale Data Streams

Kun Liu, Ruizhi Zhang & Yajun Mei

H. Milton Stewart School of Industrial and Systems Engineering

Georgia Institute of Technology

Atlanta, GA 30329-0205, USA

September 9, 2015

Abstract

In this article, motivated by biosurveillance and censoring sensor networks, we investigate the problem of distributed monitoring large-scale data streams where an undesired event may occur at some unknown time and affect only a few unknown data streams. We propose to develop scalable global monitoring schemes by parallel running local detection procedures and by combining these local procedures together to make a global decision based on SUM-shrinkage techniques. Our approach is illustrated in two concrete examples: one is the nonhomogeneous case when the pre-change and post-change local distributions are given, and the other is the homogeneous case of monitoring a large number of independent $N(0, 1)$ data streams where the means of some data streams might shift to unknown positive or negative values. Numerical simulation studies demonstrate the usefulness of the proposed schemes.

Keywords: change-point, CUSUM, parallel computing, quickest detection, sensor networks.

1 Introduction

In the modern information age, one often faces the need to online monitor large-scale data streams with the aim of offering the potential for early detection of a “trigger” event. Ideally, one would like to develop a global monitoring scheme that can detect the occurring event as quickly as possible while controlling the system-wise global false alarm rate. From the statistical point of view, this is a sequential change-point detection or quickest change detection problem, which has a variety of applications such as industrial quality control, signal detection and biosurveillance. The classical version of this problem, where one monitors independent and identically distributed (iid) *univariate* or *low-dimensional multivariate* observations from

a single data stream, is a well-developed area, and many classical procedures have been developed such as the Shewhart's chart (Shewhart [30]), moving average control charts, Page's CUSUM procedure (Page [24]), Shiryaev-Roberts procedure (Shiryaev [31], Roberts [29]), window-limited procedures (Lai [13]) and scan statistics (Glaz, Naus and Wallenstein [9]). All these classical procedures not only hold attractive theoretical properties, but also are computationally simple. See, for example, Lorden [17], Pollak [25, 26], Moustakides [22], Lai [13, 14], Kulldorff [12]. For a review, see the books such as Basseville and Nikiforov [2], Poor and Hadjiladis [27], Tartakovsky, Nikiforov, and Basseville [33].

However, research is limited in the context of monitoring large-scale data streams, especially when the occurring event might affect some, but not all, local data streams. The only exception is probably Xie and Siegmund [38], but their proposed schemes are computationally heavy with large local memory requirements to store past information, and thus is computationally infeasible for online monitoring large-scale data streams over long time period. Indeed, while many classical likelihood-ratio-based quickest change detection methods can be extended from one or low dimension to high-dimension or large-scale data streams, they are generally computationally infeasible in the context of large-scale data streams. As mentioned in Breiman [3], in order for the profession of statistics to remain healthy, more algorithm-based methods should be developed. This is exactly what needs to be done in the subfield of quickest change detection or sequential change-point detection. We feel that the current main bottleneck is on the algorithm or methodology aspect, and in particular, new ideas and new approaches are needed to develop efficient *scalable* global schemes in the sense of being able to be implemented for monitoring large-scale data streams over a long period of time.

The purpose of this article is to present a general and flexible approach that can provide efficient scalable global schemes when monitoring large-scale data streams. Our research is motivated by parallel and distributed computing and networks. A motivating example is censoring sensor networks in engineering, which was introduced by Rago, Willett, and Bar-Shalom [28] and later by Appadwedula, Veeravalli, and Jones [1] and Tay, Tsitsiklis, and Win [36]. Figure 1 illustrates the general setting of a widely used configuration of censoring sensor networks, in which the data streams $X_{k,n}$'s are observed at the remote, distributed sensors, but the final decision is made at a central location, called the fusion center. The key feature of such a network is that while sensing (i.e., taking observations at the local sensors) are generally cheap and affordable, communication between remote sensors and fusion center is expensive in terms of both energy and limited bandwidth. The question then becomes how the fusion center can still monitor the system effectively under the networks resource constraints. A more concrete example is the National Syndromic Surveillance Program BioSense Platform at the Centers for Disease Control and Prevention (CDC), where the computing power and memory of any centralized server would have become limited as compared to *daily* summary data from all state and local health departments as well as many hospitals, and thus the CDC's BioSense Platform is designed to be a distributed computing system that can make a global decision.




Figure 1: General setting of a widely used configuration of censoring sensor networks.

To develop scalable schemes for distributed monitoring large-scale data streams, we propose to take advantage of parallel computing and the fact that many efficient and computationally simple local procedures are available to detect changes in local data streams. To be more specific, suppose we are monitoring a large number K of data streams, and for each local data stream, we can construct a local detection procedure based upon some local detection statistics that can be computed recursively over time n , e.g., involving $O(1)$ computations and $O(1)$ memory requirements at each time. Then our proposed methodology is to run these K local procedures in parallel before combining them into a global monitoring scheme. Hence it only requires $O(K)$ computations and $O(K)$ memory requirements at each time step when new observations are taken, thereby yielding a scalable global monitoring scheme. While the parallel local monitoring approach sounds interesting, one allegation often made is that we will lose much information at the global level if we combine local detection procedures, not raw observation themselves, to make a global decision. Indeed, two specific methods have been developed in the literature to combine local detection procedures, but both have shortcomings: a naive method is to raise an alarm at the global level whenever any local detection procedure raises a local alarm, and the other method is developed in Mei [20] to raise a global alarm when the sum of local detection statistics is too large. Both methods are known to be inefficient when only a few (unknown) subset of data streams are affected, see Mei [20] and Xie and Siegmund [38].

In this article, we demonstrate that the problem is not on the parallel local monitoring approach itself, but on how to combine the local detection procedures suitably in the scenario when only a few (unknown) subset of data streams are affected. Our key idea is to filter out those unchanging local data streams and to make a global decision based on those likely affected data streams. For that purpose, we propose SUM-shrinkage techniques to combine the local detection statistics (in the log-likelihood ratio scale) of the local

detection procedures together to make an efficient global decision. It is worth pointing out that a well-known view in the standard off-line statistical inference literature is the necessity of shrinkage for high-dimensional data in order to improve power or efficiency. Thus, from the methodology point of view, our proposed methodologies are analogous to those off-line statistical methods such as (adaptive) truncation, and soft- and hard- thresholding, see Neyman [23], Donoho and Johnstone [5], Fan and Lin [7]. Also see Candès [4] and the references there. However, our motivation here is different and our application to quickest change detection or sequential change-point detection is new.

The remainder of this article is organized as follows. In Section 2, we present our proposed “SUM-shrinkage” methodology under a general setting of monitoring large-scale independent data streams and also discuss two existing methodologies for parallel local monitoring. We exemplify our methodology in two concrete examples: Section 3 considers the censoring sensor networks when the local data streams may or may not be homogeneous but the pre-change and post-change distributions of local data streams are given, and Section 4 investigates a more complicated scenario when the post-change distributions of local data streams involve unknown parameters. In both Sections 3 and 4, numerical Monte Carlo simulation studies are conducted to illustrate the performance of our proposed methods. Section 5 includes the proofs of Theorems 3.1 and 3.2 that justify the choices of tuning/censoring parameters in censoring sensor networks.

2 Our Proposed Methodology

Let us present our proposed methodology under a general setting, and two specific examples will be given in later sections. Assume there are K data streams in a system.

$$\begin{aligned}
 \text{Data Stream 1 :} & \quad X_{1,1}, X_{1,2}, \dots \\
 \text{Data Stream 2 :} & \quad X_{2,1}, X_{2,2}, \dots \\
 & \quad \dots \quad \dots \\
 \text{Data Stream } K : & \quad X_{K,1}, X_{K,2}, \dots .
 \end{aligned} \tag{1}$$

Initially, the system is “in control”, but at some *unknown* time ν , an undesired event may occur and affect a few unknown local data streams in the sense of changing the local distributions of the $X_{k,n}$ ’s.

Here we assume that the online monitoring is conducted under the *unstructured* environment in the sense that we do not make any assumptions to relate the occurring event to the local data streams, see Tartakovsky et al. [34], Mei [20] and Xie and Siegmund [38]. Also see Lévy-Leduc and Roueff [15] for an application of the unstructured problem to anomaly detection in computer networks. In particular, we focus on the scenario when the occurring event changes the local distributions of affected local data streams, and we do not aim to detect changes on the correlation between different data streams. Hence, the data $X_{k,n}$ ’s will be assumed

to be independent across different data streams, but can be flexible otherwise. For instance, the $X_{k,n}$'s may or may not be identically distributed across different local data streams, can be dependent over time within each local data stream, and can be *univariate* or *low-dimensional multivariate*. We should mention that the assumption of the independence across different data streams is standard in the unstructured problem, see Tartakovsky et al. [34], Mei [20] and Xie and Siegmund [38]. This is not as restrictive as one thought in many practical applications, as the $X_{k,n}$'s can be chosen as the residuals of some spatio-temporal models rather than the original raw data. In other words, one can first pre-process the data, build a spatio-temporal baseline model when the system is in control, calculate the corresponding residuals, and then monitor the residuals $X_{k,n}$'s. For an illustration, see Xie, Huang and Willett [37], and Liu, Mei and Shi [16], to monitor the residuals of dependent data in two real-world applications in solar flare and hot-forming process.

For the purpose of generalization, we do not specify which kind of local changes these K data streams may have. Instead we assume that there is a local detection statistic $W_{k,n}$ (in the log-likelihood scale) for the k -th local data stream at each time step n that summarizes the evidence regarding a possible local change based on the first n local observations $(X_{k,1}, \dots, X_{k,n})$ for each $k = 1, \dots, K$. It is important that the $W_{k,n}$'s not only can detect local changes quickly, depending on specific assumptions on the local changes and local data distribution model, but also can be computed recursively. We should emphasize that it can be highly non-trivial to construct such $W_{k,n}$'s, especially when the local post-change distributions involve unknown parameters or when missing data are present, see our example in Section 4 and Liu, Mei and Shi [16]. Here we assume, for a moment, that such $W_{k,n}$'s have been constructed, and our focus is the parallel local monitoring method that combines these local detection statistics $W_{k,n}$'s together to make an efficient global decision.

Before presenting our proposed methodology, let us first review the definition of global false alarm rate and two existing methods for parallel local monitoring. When monitoring K independent data streams in (1), it is well-known in statistics that even if each local false alarm rate is well controlled, the global false alarm rate can be significant when the number K of data streams is large. In the literature of sequential change-point detection, for a global monitoring scheme that raise an alarm at time T , its global false alarm rate is often evaluated by $1/\mathbf{E}^{(\infty)}(T)$, where $\mathbf{E}^{(\infty)}(T)$ is the expectation of T when the system is "in control," and is often called the average run length to false alarm. A standard global false alarm constraint is to require a global monitoring scheme with a stopping time T satisfying

$$\mathbf{E}^{(\infty)}(T) \geq \gamma, \quad (2)$$

where $\gamma > 0$ is a pre-specified constant. The rigorous definition of the detection delay of the scheme T will be postponed later in next section.

There are two existing methods for parallel local monitoring. The first one is to raise an alarm at the global level whenever any local detection procedures raises a local alarm. If we normalize the local detection

statistics $W_{k,n}$'s, this can be rewritten as raising an alarm at the global level at time

$$T_{\max}(a) = \inf\{n \geq 1 : \max_{1 \leq k \leq K} W_{k,n} \geq a\}, \quad (3)$$

($= \infty$ if such n does not exist) where $a > 0$ is a pre-specified constant. Below we will call the scheme in (3) the “MAX” scheme. The second method is the “SUM” scheme developed in Mei [20] that is defined by the stopping time

$$T_{\text{sum}}(a) = \inf\{n \geq 1 : \sum_{k=1}^K W_{k,n} \geq a\}, \quad (4)$$

($= \infty$ if such n does not exist). As mentioned in Mei [20], the “MAX” scheme $T_{\max}(a)$ in (3) works well when one or very few data streams are affected, whereas the “SUM” scheme $T_{\text{sum}}(a)$ in (4) works well only when many data streams are affected. Here and below the threshold a of a scheme $T(a)$ is a pre-specified constant so that the scheme $T(a)$ satisfies the false alarm constraint γ in (2).

Now we are ready to present our proposed methodology under a general setting. We suggest to define the global monitoring statistic of the general “SUM-shrinkage” form

$$G_n = \sum_{k=1}^K h_k(W_{k,n}), \quad (5)$$

where $h_k(\cdot) \geq 0$ are some suitable shrinkage transformation functions. Then our proposed SUM-shrinkage scheme raises a global alarm at the time

$$N_G(a) = \inf\{n \geq 1 : G_n \geq a\}. \quad (6)$$

Intuitively, the shrinkage functions h_k 's in (5) play the role of dimension reduction by automatically filtering out those non-changing local data streams and by focusing only on those local data streams that appear to be affected by the occurring event. In addition, as an extension, the proposed SUM-shrinkage scheme $N_G(a)$ enjoys the nice properties of the SUM scheme $T_{\text{sum}}(a)$ in (4): it does not assume that all local data streams are affected by the occurring event simultaneously, and thus can be very useful when there may be a time delay before the event affects different local data streams, or when different local data streams involve local changes at different time steps. Also see Mei [20] and Xie and Siegmund [38] for more discussions.

Evidently a suitable choice of the h_k 's in the SUM-shrinkage monitoring statistic G_n in (5) will depend on the assumptions and contexts of applications. In Sections 3 and 4 below we will demonstrate the following three shrinkage transformations of the form

- Hard-thresholding: $h(x) = x \mathbf{1}\{x \geq b\}$ for some constant b , (7)

- Soft-thresholding: $h(x) = \max\{x - b, 0\}$ for some constant b , (8)

- Order-thresholding: $h(x) = x \mathbf{1}\{x \geq w_{(r)}\}$, where $w_{(r)}$ is the r -th largest statistic of w_1, \dots, w_K . (9)

The rationale and motivations of these three transformations will be given in Sections 3 and 4. Of course, besides those in (7)-(9), there are many other kinds of the shrinkage functions such as $h(x) = \exp(bx)$. Also by semi-Bayesian arguments, the transformation $h(x) = \log[1 - p_0 + p_0 \exp(x)]$ is proposed and used in the schemes of Xie and Siegmund [38] in a completely different manner under their setting.

We should emphasize that our proposed SUM-shrinkage schemes can be easily implemented in the distributed network systems as long as the local detection statistics $W_{k,n}$'s can be easily computed at local data streams. Besides the two examples in the next two sections, we should point out that our proposed "SUM-shrinkage" methodology in (5)-(6) has a broad range of other applications. For instance, the local detection statistics $W_{k,n}$'s can be defined for dependent observations such as those from the recursive schemes in Fuh and Mei [8] for hidden Markov models, or those from the non-parametric detection schemes in Gorden and Pollak [10], depending on which kind of local models or local changes we are interested in. In addition, little information seems to be lost if we do not observe those local data streams with small values of $W_{k,n}$'s since they make limited contributions in our proposed global monitoring statistic G_n in (5). This motivated Liu, Mei and Shi [16] to develop an efficient adaptive sensor relocation policy when one only has ability to observe r out of K data streams at each time step. This may occur in manufacturing process control when there are K possible stages in the process but there are only r expensive sensors available to monitor the process. In such a problem, the order-thresholding transformation can be combined with missing data techniques to be used not only in the global monitoring statistic G_n in (5) for quickest detection, but also in a greedy manner to adaptively observe those r data streams with the largest $W_{k,n}$'s values at each time step. We feel the spirit of SUM-Shrinkage can have many other applications, and hopefully our research opens new research opportunities and directions, especially on monitoring large-scale data streams.

2.1 General Guidelines

Below we will provide some general guidelines on how to use our proposed SUM-shrinkage scheme $N_G(a)$ in (6). The performance of our proposed scheme will depend on two components: one is the local detection statistics $W_{k,n}$'s and the other is the shrinkage transformation functions h_k 's. The suitable choices of these two components depend on the applications and contexts, and will be demonstrated in the next two sections.

In general, the local detection statistics $W_{k,n}$'s should be able to efficiently detect local changes we are interested in, and ideally can also be recursively computed over time so that the local computation is simple. Such a choice of the $W_{k,n}$'s might be straightforward in some applications (i.e., Section 3), but can be highly non-trivial in other cases (e.g., Section 4). Also see Liu, Mei and Shi [16] how to define $W_{k,n}$'s when the observations from some data streams are missing or unobservable.

For the choice of the shrinkage transformation $h_k(\cdot)$'s, the situation can be complicated when the data streams are nonhomogeneous (or the $W_{k,n}$'s have different properties for different k 's). One rule of thumb

is to choose $h_k(\cdot)$'s so that the local procedures based upon nonhomogeneous $W_{k,n}$'s will raise local alarms at roughly same time steps for all affected data streams, see our example in Section 3.

Given the choices of the local detection statistics $W_{k,n}$'s and the shrinkage transformation $h_k(\cdot)$'s, an important remaining question is how to determine the global threshold a in (6) so that the proposed SUM-shrinkage scheme $N_G(a)$ in (6) satisfies the global false alarm constraint γ in (2). In general this is nontrivial, as it requires one to accurately characterize the relationship between the threshold a and the false alarm constraint γ , when the dimension K goes to ∞ . Intuitively, the global monitoring statistic G_n in (5) is the sum of K (independent) random variables, one would expect that the central limited theorem (CLT) will be useful when the shrinkage transformation keeps most non-zero values, e.g., the hard-thresholding or soft-thresholding transformations in (7) or (8) when the censoring parameters b 's are not large, whereas the compound Poisson process will be needed when the shrinkage transformation only keeps very few non-zero values, e.g., the order-thresholding transformation in (9) with not so large r value. The rigorous theoretical proofs are beyond the scope of this article and will be investigated elsewhere. Below we will use Chebyshev's inequality and CLT to provide two approximations of the global threshold a in terms of γ .

To do so, let us assume that under the pre-change hypothesis $\mathbf{P}^{(\infty)}$, the local detection statistics $W_{k,n}$'s and their shrinkage transformations converge very fast to stationary distributions as the time step n goes to ∞ . More specifically, we assume that for each k , the $h_k(W_{k,n})$'s converge to their limit H_k^* which is stochastically larger than $h_k(W_{k,n})$'s and has a well-defined log-moment generating function

$$\psi_k(\theta) = \log \mathbf{E}^{(\infty)} \exp(\theta H_k^*)$$

for all $\theta \geq 0$. By the definition of $N_G(a)$ in (6) and by applying Chebyshev's inequality to both $N_G(a) \geq 0$ and $\sum_{k=1}^K H_k^*$, for any $x > 0$,

$$\begin{aligned} \mathbf{E}^{(\infty)}(N_G(a)) &\geq x \mathbf{P}^{(\infty)}(N_G(a) \geq x) \\ &= x \left[1 - \mathbf{P}^{(\infty)}(N_G(a) < x) \right] \\ &= x \left[1 - \mathbf{P}^{(\infty)}\left(\sum_{k=1}^K h_k(W_{k,n}) \geq a \text{ for some } 1 \leq n \leq x\right) \right] \\ &\geq x \left[1 - x \mathbf{P}^{(\infty)}\left(\sum_{k=1}^K H_k^* \geq a\right) \right] \\ &\geq x \left[1 - x e^{-\theta a} \mathbf{E}^{(\infty)} \exp\left(\theta \sum_{k=1}^K H_k^*\right) \right] \\ &= x \left[1 - x e^{-\theta a} \exp\left(\sum_{k=1}^K \psi_k(\theta)\right) \right], \end{aligned}$$

where the last equation uses the fact that these K data streams are independent across different data streams.

Define

$$u_\theta(a) = e^{-\theta a} \exp\left(\sum_{k=1}^K \psi_k(\theta)\right).$$

Choosing x to maximize $x(1 - xu_\theta(a))$, we have $x = 1/(2u_\theta(a))$, and thus $\mathbf{E}^{(\infty)}(N_G(a)) \geq 1/(4u_\theta(a))$ for all $a > 0$ and all $\theta > 0$. Thus, in the special case when there exists a $\theta_0 > 0$ so that $\psi_k(\theta_0) \leq 0$ for all k , we have $u_\theta(a) \leq \exp(-\theta_0 a)$ and thus a simple choice of

$$a = (\log \gamma + \log 4)/\theta_0 \tag{10}$$

will guarantee that $\mathbf{E}^{(\infty)}(N_G(a)) \geq \gamma$. This implies that a is of order $(\log \gamma)/\theta_0$, where the θ_0 value depends on the shrinkage transformation $h_k(\cdot)$.

Unfortunately, our numerical simulation suggests that the above choice of a based on Chebyshev's inequalities is often too loose. A better estimation of the global threshold a can be found heuristically by using more refined approximations. When the CLT is applicable to the global monitoring statistic G_n in (5), we can use the approximation

$$\mathbf{P}^{(\infty)}\left(\sum_{k=1}^K H_k^* \geq a\right) \approx \mathbf{P}\left(N(0, 1) \geq \frac{a - \mu_H}{\sigma_H}\right),$$

where μ_H and σ_H^2 are the mean and variance of the limiting global statistic $\sum_{k=1}^K H_k^*$:

$$\mu_H = \sum_{k=1}^K \dot{\psi}_k(0) \quad \text{and} \quad \sigma_H^2 = \sum_{k=1}^K \ddot{\psi}_k(0).$$

In addition, if we approximate the distribution of $N_G(a)$ as an exponentially distribution, which is true to most sequential change-point detection schemes in the literature, then we have $\mathbf{E}^{(\infty)}(N_G) \approx x/\mathbf{P}^{(\infty)}(N_G(a) \leq x)$ for moderately large x . Combining these above two approximations yields a heuristic approximation

$$a = \mu_H + z_{1/\gamma}\sigma_H, \tag{11}$$

where $z_{1/\gamma} = z$ so that $\mathbf{P}(N(0, 1) \geq z) = 1/\gamma$. Our numerical simulation below supports this heuristic approximation, but rigorous justifications turn out to be highly technical due to the complicated correlation structures of the $W_{k,n}$'s over time domain n , and will be presented elsewhere.

3 A First Example: Censoring Sensor Networks with Known Post-Change Distributions

For censoring sensor networks in Figure 1, practitioners often prefer the local sensors to send summary messages $U_{k,n}$'s to the fusion center only when necessary, so as to prolong the reliability and lifetime of

the network system. The question then becomes when and how to send summary messages so that the fusion center can still monitor the network system effectively. In the quickest change detection problem in censoring sensor networks, let $X_{k,n}$ denote the observations at the k -th sensor at time step n . In this section, we focus on nonhomogeneous sensors, but make a restrictive assumption that the pre-change and post-change distributions of the $X_{k,n}$'s are given. In the next section, we will investigate the case of homogeneous sensors when the post-change distributions are unknown.

Specifically, in this section we follow the literature to assume that for each $k = 1, \dots, K$, the density function of the observations at the k -th data stream is f_k before the change, and is g_k after the change if the k -th data stream is affected, where the f_k 's and g_k 's are completely specified densities with respect to a suitable measure μ , see, for example, Tartakovsky and Veeravalli [35]. For each $1 \leq k \leq K$, we assume that the Kullback-Leibler (KL) information number

$$I(g_k, f_k) = \int \log \frac{g_k(x)}{f_k(x)} g_k(x) d\mu(x) \quad (12)$$

is finite and positive, and

$$\int \left(\log \frac{g_k(x)}{f_k(x)} \right)^2 g_k(x) d\mu(x) < \infty. \quad (13)$$

3.1 Our Proposed Methods

Let us apply our proposed SUM-shrinkage schemes to censoring sensor networks in Figure 1. To do so, we need to define two components of our proposed global monitoring statistics in (5). The first one is the local detection statistic $W_{k,n}$'s, which is simple in this context since the pre-change and post-change distributions, f_k and g_k , are known at each local sensor. For instance, the $W_{k,n}$'s can be chosen as the well-known local CUSUM statistics (Page [24]) that are defined recursively by

$$W_{k,n} = \max \left(W_{k,n-1} + \log \frac{g_k(X_{k,n})}{f_k(X_{k,n})}, 0 \right), \quad (14)$$

for $n \geq 1$ and $W_{k,0} = 0$ for $k = 1, \dots, K$. As shown in Lorden [17] and Moustakides [22], the local CUSUM statistics $W_{k,n}$'s in (14) yield the optimal local procedure to detect the local change under some suitable criteria.

Next, we need to specify concrete shrinkage transformation h_k 's in (5) for censoring sensor networks. To prolong the reliability and lifetime of the network system, it is natural for the local sensors to transmit only those local CUSUM statistics $W_{k,n}$'s that are large. Specifically, at time n , the sensor message from the sensor to the fusion center is given by

$$U_{k,n} = \begin{cases} W_{k,n}, & \text{if } W_{k,n} \geq b_k \\ \text{NULL}, & \text{if } W_{k,n} < b_k \end{cases}, \quad (15)$$

where $b_k \geq 0$ is the local censoring parameter at the k -th sensor (or data stream). In practice, the message “NULL” could be represented by the situation when the sensor does not send any message to the fusion center, e.g., the sensor is silent.

After receiving the local sensor messages from the sensors, the fusion center then combines these local sensor messages $U_{k,n}$ ’s in (15) suitably together to make a global decision. There are many approaches to do so, and below we illustrate three of them. The first two schemes are based on the summation of all sensor messages $U_{k,n}$ ’s, depending on how to interpret the “NULL” values. If we treat the “NULL” values as lower limit 0, then the fusion center raises a global alarm at time

$$\begin{aligned} N_{hard}(a) &= \inf \left\{ n \geq 1 : \sum_{k=1}^K U_{k,n} \geq a \right\} \\ &= \inf \left\{ n \geq 1 : \sum_{k=1}^K W_{k,n} \mathbf{1}\{W_{k,n} \geq b_k\} \geq a \right\}. \end{aligned} \quad (16)$$

Below this scheme will be referred as the hard-thresholding scheme, since it is a special case of the global statistic in (5) when the shrinkage functions h_k ’s are the hard-thresholding transformation in (7).

Meanwhile, if we treat the “NULL” values as the upper limit b_k ’s, then the fusion center will compute the global monitoring statistic

$$G_n = \sum_{k=1}^K U_{k,n} = \sum_{k=1}^K \max\{W_{k,n}, b_k\} = \sum_{k=1}^K \max\{W_{k,n} - b_k, 0\} + \sum_{k=1}^K b_k,$$

which is closely related to the soft-thresholding transformation in (8). Hence, we can define the soft-thresholding scheme that raises an alarm at time

$$N_{soft}(a) = \inf \left\{ n \geq 1 : \sum_{k=1}^K \max\{W_{k,n} - b_k, 0\} \geq a \right\}. \quad (17)$$

Here we keep the threshold of $N_{soft}(a)$ as a instead of $a - \sum_{k=1}^K b_k$, so that $N_{soft}(a)$ is the special case of our proposed SUM-shrinkage scheme $N_G(a)$ in (6) with the soft-thresholding transformation in (8).

The third approach occurs when the fusion center has a prior knowledge that (at most) r out of K data streams will be affected by the occurring event. Such a prior knowledge may be defined by the network fault-tolerant design to avoid risking failure. In this case, it is reasonable for the fusion center to order all sensor messages $U_{k,n}$ ’s as $U_{(1),n} \geq \dots \geq U_{(K),n}$, and raise an alarm if the sum of the r largest $U_{k,n}$ ’s is too large. This is a combination of the hard-thresholding transformation in (7) and the order-thresholding transformation in (9), and it yields a global scheme that is defined by the stopping time

$$N_{comb,r}(a) = \inf \left\{ n \geq 1 : \sum_{k=1}^r U_{(k),n} \geq a \right\}. \quad (18)$$

For simplicity, the “NULL” values of $U_{k,n}$ ’s in the scheme $N_{comb,r}(a)$ in (18) will be treated as the lower limit 0 in our simulation below.

For the purpose of comparison, we also apply the order-thresholding transformation in (9) directly to the local CUSUM statistics $W_{k,n}$'s in (14) themselves. Specifically, we order the K local CUSUM statistics $W_{1,n}, \dots, W_{K,n}$ from largest to smallest: $W_{(1),n} \geq W_{(2),n} \geq \dots \geq W_{(K),n}$. Then the order-thresholding scheme can be defined by the stopping time

$$N_{order,r}(a) = \inf \left\{ n \geq 1 : \sum_{k=1}^r W_{(k),n} \geq a \right\}. \quad (19)$$

Of course, $N_{order,r}(a)$ is a special case of $N_{comb,r}(a)$ when the local censoring parameter $b_k \equiv 0$, since the local CUSUM statistics $W_{k,n} \geq 0$ for all k and all n .

It is useful to mention that each of all four schemes in (16)-(19) is based on our proposed shrinkage statistics in (5), but each is actually a very large family of schemes that includes “MAX” or “SUM” schemes or both as special cases. For instance, for the hard-thresholding scheme $N_{hard}(a)$ in (16), it becomes the “SUM” scheme $T_{sum}(a)$ in (4) if the censoring parameter $b_k \equiv 0$ for all k , but becomes the “MAX” scheme $T_{max}(a)$ in (3) if $b_k \equiv a$ for all k . Similarly, the order-thresholding scheme $N_{order,r}(a)$ in (19) becomes the “MAX” scheme when the order parameter $r = 1$ and becomes the “SUM” scheme when $r = K$. As for the soft-thresholding scheme $N_{soft}(a)$ in (17), it becomes the “SUM” scheme if $b_k \equiv 0$ for all k , and based on our numerical experience, its properties are similar to those of the “MAX” scheme when b_k 's are very large.

Besides the local CUSUM statistics, another popular local detection statistic is the local Shiryaev-Roberts statistic (Shiryaev [31], Roberts [29]) which can be defined in the log-likelihood ratio scale by

$$\hat{W}_{k,n} = \log \left(\exp(\hat{W}_{k,n-1}) + 1 \right) + \log \frac{g_k(X_{k,n})}{f_k(X_{k,n})} \quad (20)$$

for $n \geq 1$ and $\hat{W}_{k,0} = 0$. It is well-known that the local Shiryaev-Roberts statistics $\hat{W}_{k,n}$ in (20) yield an efficient local detection procedure whose performance is similar to that of local CUSUM statistics in (14) when detecting a local change in distribution from f_k to g_k , see Pollak [25, 26]. In our numerical analysis below, the local detection statistics $W_{k,n}$'s can also be defined as $\hat{W}_{k,n}$'s in (20), the local Shiryaev-Roberts statistics in logarithm scale, or better yet, its positive part $\max\{\hat{W}_{k,n}, 0\}$. Our numerical simulation experiences suggest that the performances of global monitoring schemes based upon local Shiryaev-Roberts statistics are similar to those based upon local CUSUM statistics in (14) when monitoring K data streams. Unfortunately it is still an open question to investigate the theoretical properties of Shiryaev-Roberts-type schemes in the context of K data streams, and thus we will focus on the local CUSUM statistics $W_{k,n}$'s in (14) as the local detection statistics below.

3.2 Choices of the Thresholds b_k 's

So far we simply follow our intuition without discussing how to choose the local censoring parameters b_k 's in (15) for censoring sensor networks. Intuitively, the b_k 's should be the same when the sensors are homogeneous,

but they probably should be different when the sensors are nonhomogeneous. It turns out that a “good” choice is

$$b_k = \rho_k b \quad (21)$$

for $k = 1, \dots, K$ for some common constant $b \geq 0$, where

$$\rho_k = \frac{I(g_k, f_k)}{\sum_{k=1}^K I(g_k, f_k)} \quad (22)$$

and $I(g_k, f_k)$ is the KL information number defined in (12). Theoretical justification of our choice of b_k in (21)-(22) will be postponed to subsection 3.4. Roughly speaking, ρ_k in (22) can be thought of as the weight of the k -th data stream in the overall final decision, and the choice of $b_k = \rho_k b$ in (21) allows those affected local sensors to send local messages $U_{k,n}$ ’s with large values to the fusion center at roughly the same time, thereby leading the quick detection of occurring event.

It remains to choose the common constant $b > 0$ in (21). This may be determined by a non-statistical constraint in censoring sensor networks that the average fraction of transmitting sensors at any time step is restricted to be at most $\eta \in (0, 1)$ when no change occurs. In this case, when no event occurs, the average fraction of transmitting sensors at any time step n is

$$\frac{1}{K} \sum_{k=1}^K \mathbf{P}^{(\infty)}(U_{k,n} \neq \text{NULL}) = \frac{1}{K} \sum_{k=1}^K \mathbf{P}^{(\infty)}(W_{k,n} \geq \rho_k b) \leq \frac{1}{K} \sum_{k=1}^K \exp(-\rho_k b) \leq \exp(-\rho_{\min} b),$$

where $\rho_{\min} = \min_{1 \leq k \leq K} \rho_k$ and the second-to-last inequality follows from the well-known properties of the local CUSUM statistics that $\mathbf{P}^{(\infty)}(W_{k,n} \geq a) \leq \exp(-a)$ for all $a > 0$, see, for example, Appendix 2 on Page 245 of Siegmund [32]. Thus a choice of $b = (1/\rho_{\min}) \log \eta^{-1}$ will guarantee that on average, at most $100\eta\%$ of K sensors will transmit messages at any given time when no event occurs.

A special case occurs when all K sensors are homogeneous in the sense that the KL information numbers $I(g_k, f_k)$ ’s in (12) are the same for all k . Then we have $\rho_{\min} = 1/K$, and our proposed choice of the local censoring parameter is given by

$$b_k = \rho_k b = (1/K)(K \log \eta^{-1}) = \log \eta^{-1}, \quad (23)$$

for all $k = 1, \dots, K$. It is interesting to see that as the number K of homogeneous sensors increases, the weight ρ_k of each local sensor in the overall final decision is decreasing, but the common constant b is increasing. Thus the choice of the local censoring parameter b_k ’s in (23) remains as a constant, and this seems attractive to practitioners in censoring sensor networks.

It is important to emphasize for each of our proposed schemes $T(a)$ in (16)-(19), the stopping time $T(a)$ is increasing as a function of the censoring parameters b_k ’s when the global threshold value a is *given*. That is, a larger value of b_k ’s implies both larger ARL to false alarm and larger detection delays. However, the

situation becomes completely different when $T(a)$ is required to satisfy the false alarm constraint (2). This is because different global threshold values a 's are needed for these schemes with different b_k 's, and thus larger values of b_k 's may or may not lead to larger detection delays. Also see our numerical simulations below.

3.3 Numerical Simulations

In this subsection we report our numerical simulation results to illustrate the usefulness of the proposed schemes in (16)-(19). Suppose that there are $K = 100$ independent and identical sensors in a system, and the observations at each sensor are iid with mean 0 and variance 1 before the change and with mean 1 and variance 1 after the change if affected. In our simulation study, we simply assume that the change is instantaneous if a sensor is affected, but we do not know which subset of sensors will be affected by the occurring event.

For the purpose of comparison, we conduct numerical simulations for six families of global monitoring schemes:

- the “MAX” scheme $T_{\max}(a)$ in (3),
- the “SUM” scheme $T_{\text{sum}}(a)$ in (4),
- the order thresholding scheme $N_{\text{order},r}(a)$ in (19) with $r = 10$,
- the hard thresholding scheme $N_{\text{hard}}(a)$ in (16),
- the soft thresholding scheme $N_{\text{soft}}(a)$ in (17),
- the combined thresholding schemes $N_{\text{comb},r}(a)$ in (18) with $r = 10$.

The first three schemes require all local sensors to send all local CUSUM statistics $W_{k,n}$'s values to the fusion center at each and every time step, and corresponds to the case when the local censoring parameter $b_k \equiv 0$ for all $k = 1, \dots, K$. For order-thresholding in the families of $N_{\text{order},r}(a)$ and $N_{\text{comb},r}(a)$, we choose $r = 10$ to better understand the scenario when 10 out of 100 sensors are affected by the occurring event. For each of the last three schemes in the list, i.e., our three proposed schemes (16)-(18), we further consider three different values of the local censoring parameters b_k 's:

- (i) $b_k \equiv 1/2 \approx -\log(0.607)$ for all k ,
- (ii) $b_k \equiv -\log(0.1) = 2.3026$ for all k ,
- (iii) $b_k \equiv -\log(0.01) = 4.6052$ for all k .

The choices of these values will guarantee that when no event occurs, on average at most $\eta = 60.7\%$, 10% , and 1% of $K = 100$ homogeneous sensors will transmit messages at any given time, respectively. Therefore, there are a total of $3 + 3 * 3 = 12$ specific schemes in our numerical simulation study.

For each of these 12 specific schemes $T(a)$, we first find the appropriate values of the global threshold a to satisfy the false alarm constraint $\mathbf{E}^{(\infty)}(T(a)) \approx \gamma = 5000$ (within the range of sampling error). Next, using the obtained global threshold value a , we simulate the detection delay when the change-point occurs at time $\nu = 1$ under several different post-change scenarios, i.e., different number of affected sensors. All Monte Carlo simulations are based on $m = 2500$ repetitions.

Table 1 summarizes our simulated detection delays of these 12 schemes under 8 different post-change hypothesis, depending on the number of affected sensors. From Table 1, among these 12 specific schemes, when a small number ($1 \sim 3$) of 100 homogeneous sensors are affected by the event, the “MAX” scheme $T_{\max}(a)$ is the best (in the sense of smallest detection delay), the “SUM” scheme $T_{\text{sum}}(a)$ is the worst, and all other schemes are in-between. Similarly, when a large number (20 or more) of 100 homogeneous sensors are affected, the order is reserved: $T_{\text{sum}}(a)$ is the best, $T_{\max}(a)$ is the worst, and all other schemes are in-between. However, when $5 \sim 10$ sensors are affected, the schemes with order-thresholding $r = 10$ yield the smallest detection delays, since they are designed to detect the scenario when 10 sensors are affected by the event. An interesting observation is that the soft-thresholding scheme $N_{\text{soft}}(a)$ can also yield the smallest detection delays with a suitable choice of b_k ’s. In addition, it is clear from Table 1 that for each given scheme, the fewer affected sensors we have, the larger detection delay it will have. All these results are consistent with our intuition.

It is worth emphasizing that for the families of the hard-thresholding schemes $N_{\text{hard}}(a)$ in (16) or the soft-thresholding schemes $N_{\text{soft}}(a)$ in (17), a larger censoring value of b_k actually leads to a smaller detection delay when only a few sensors (between 1 and 5 sensors) are affected. This suggests that a larger censoring value b_k may actually be necessary for efficient detection when the affected sensors are sparse.

A surprising and possibly counter-intuitive result in Table 1 is the effect of not so large values of censoring parameters b_k ’s in finite sample simulations. For instance, the performances of the “SUM” scheme $T_{\text{sum}}(a)$ and the hard thresholding scheme $N_{\text{hard}}(a, b_k = 0.50)$ are similar in view of sampling errors. Likewise, the top- r thresholding scheme $N_{\text{order}, r=10}(a)$ and the combined thresholding scheme $N_{\text{comb}, r=10}(a, b_k = 0.50)$ also have identical performances. The interpretation in the censoring sensor networks context is as follows: using our proposed communication policy in (15), we only need $\exp(-b_k) = \exp(-0.5) = 60.7\%$ of 100 sensors to transmit information to the fusion center at any given time when no event occurs, but we can still be as effective as the full transmission scenario when all sensors transmit information at all time steps. In other words, much communication costs can be saved by our proposed schemes $N_{\text{hard}}(a)$ or $N_{\text{comb}, r}(a)$ with not so large values of b_k ’s.

Table 1: A comparison of the detection delays of six families of schemes with $\gamma = 5000$. The smallest and largest standard errors of these 12 schemes are also reported under each post-change hypothesis based on 2500 repetitions in Monte Carlo simulations.

	# sensors affected								
	1	3	5	8	10	20	30	50	100
Smallest standard error	0.18	0.07	0.05	0.03	0.03	0.02	0.01	0.01	0.00
Largest standard error	0.35	0.12	0.07	0.06	0.05	0.04	0.03	0.03	0.03
Schemes with $b_k \equiv 0$									
$T_{\max}(a = 11.27)$	23.3	16.3	14.4	13.0	12.4	10.9	10.2	9.5	8.7
$T_{\text{sum}}(a = 88.66)$	52.1	21.8	14.7	10.3	8.7	5.2	3.9	2.9	2.0
$N_{\text{order}, r=10}(a = 44.11)$	34.1	15.5	11.2	8.5	7.5	5.5	4.8	4.1	3.4
Schemes $N_{\text{hard}}(a)$ in (16) with different positive b_k 's									
$N_{\text{hard}}(a = 85.60, b_k = 0.50)$	52.9	21.9	14.9	10.3	8.7	5.2	4.0	2.9	2.0
$N_{\text{hard}}(a = 52.21, b_k = 2.3026)$	50.6	20.7	13.8	9.6	8.2	5.2	4.2	3.2	2.4
$N_{\text{hard}}(a = 26.31, b_k = 4.6052)$	39.8	16.0	11.5	8.8	7.9	5.9	5.2	4.4	3.8
Schemes $N_{\text{soft}}(a)$ in (17) with different positive b_k 's									
$N_{\text{soft}}(a = 63.92, b_k = 0.50)$	48.2	20.2	13.7	9.7	8.2	5.1	4.0	3.0	2.0
$N_{\text{soft}}(a = 21.56, b_k = 2.3026)$	33.9	15.4	11.2	8.5	7.5	5.3	4.5	3.7	3.0
$N_{\text{soft}}(a = 8.29, b_k = 4.6052)$	25.2	13.8	11.1	9.2	8.4	6.7	5.9	5.2	4.4
Schemes $N_{\text{comb}, r}(a)$ in (18) with $r = 10$ and different positive b_k 's									
$N_{\text{comb}, r}(a = 44.11, b_k = 0.50)$	34.1	15.5	11.2	8.5	7.5	5.5	4.8	4.1	3.4
$N_{\text{comb}, r}(a = 43.88, b_k = 2.3026)$	38.5	16.8	11.7	8.6	7.5	5.5	4.7	4.0	3.3
$N_{\text{comb}, r}(a = 26.31, b_k = 4.6052)$	39.8	16.0	11.5	8.8	7.9	5.9	5.2	4.4	3.8

It is also interesting to see the effect of the order-thresholding parameter r in finite sample simulations when the hard-thresholding parameters b_k 's are large. From Table 1, when the false alarm constraint γ in (2) is only moderately large, e.g., $\gamma = 5000$, the performances of $N_{hard}(a, b_k)$ and $N_{comb, r=10}(a, b_k)$ are identical when $b_k = 4.6052$ — they not only have the same global threshold a , but also have the same detection delays. Intuitively, the stopping time $N_{comb, r}(a, b_k)$ is decreasing as a function of r , and thus we have $N_{hard}(a, b_k) = N_{comb, r=K}(a, b_k) \leq N_{comb, r=10}(a, b_k)$ when $b_k = 4.6052$. So one may wonder why our numerical simulations lead to identical results? One explanation is that with such a choice of $b_k = 4.6052$, when no event occurs, on average there is at most 1 non-zero sensor message received in the fusion center at any given time, and thus there is little difference whether one uses the sum of the largest $r = 10$ sensor messages or

uses the sum of all $K = 100$ sensor messages. Hence similar performances are observed in finite-sample simulations.

3.4 Asymptotic Optimality Theory

In this subsection, we provide theoretical justification of our choices of the local censoring parameters b_k 's in (21)-(22) and we will show the corresponding schemes hold certain asymptotic optimality properties. To emphasize the choices of $b_k = \rho_k b$ in (21)-(22) with $b \geq 0$ being the common constant, we rewrite our proposed schemes as $N_{hard}(a, b)$, $N_{soft}(a, b)$ and $N_{comb, r}(a, b)$ in this subsection and only in this subsection.

Let us begin with a rigorous definition of the post-change hypothesis. We assume that the k -th data stream is affected at time $\nu_k = \nu + \delta_k$, where the term $\delta_k \in [0, \infty]$ denotes the delay of the occurring event's impact on the k -th data stream, and $\delta_k = \infty$ implies that the k -th data stream is not affected. That is, the density function of the sensor observations $X_{k, n}$'s of the k -th data stream changes from f_k to g_k at time $\nu_k = \nu + \delta_k$. In the case when the change is instantaneous, the delay effect δ_k only takes two possible values, 0 or ∞ . Here we relax such an assumption a little bit, and assume that the change might not be instantaneous.

To simplify our arguments and highlight our main ideas, we will assume that the delay effects δ_k 's satisfy the following post-change hypothesis set Δ :

$$\Delta = \{(\delta_1, \dots, \delta_K) : \text{the } \delta_k \text{'s either } = \infty \text{ or satisfy } 0 \leq \delta_k < \log \gamma \text{ and } \min_{1 \leq k \leq K} \delta_k = 0\}. \quad (24)$$

where γ is the false alarm constraint in (2), and $x(t) < y(t)$ implies that $x(t)/y(t) \rightarrow 0$ as $t \rightarrow \infty$. Note that the assumption of $\min_{1 \leq k \leq K} \delta_k = 0$ is trivial, since otherwise the system is actually affected by the occurring event at the “new” change-point $\nu' = \nu + \min_{1 \leq k \leq K} \delta_k$. The assumption of $\delta_k < \log \gamma$ is a technical assumption to ensure that one is able to utilize all affected data streams to raise a global alarm subject to the false alarm constraint γ in (2). In other words, we only consider the scenario when the differences on the finite delay affects δ_k 's are not too large as compared to the typical order $(\log \gamma)$ of detection delays. A

sufficient condition to satisfy this assumption is when all finite δ_k 's are uniformly bounded by some constants that do not depend on the false alarm constraint γ in (2).

Next, let us define the detection delay of a global monitoring scheme rigorously when the event occurs at the unknown time ν with specific delay effects δ_k 's. Suppose a global monitoring scheme raises an alarm at time $T \geq \nu$, it takes $T - \nu + 1$ time steps from the post-change scenario to indicate that an event might occur, and thus $T - \nu + 1$ can be regarded as the detection delay. To take into account of the randomness of T and the uncertainty of ν , a widely used rigorous definition of the detection delay of T is the following "worst case" detection delay defined in Lorden [17],

$$\bar{\mathbf{E}}_{\delta_1, \dots, \delta_K}(T) = \sup_{\nu \geq 1} \text{ess sup } \mathbf{E}^{(\nu)} \left((T - \nu + 1)^+ \middle| \mathcal{F}_{\nu-1} \right).$$

Here the δ_k 's are the delay effects, $\mathcal{F}_{\nu-1} = (X_{1,[1,\nu-1]}, \dots, X_{K,[1,\nu-1]})$ denotes past global information at time ν , $X_{k,[1,\nu-1]} = (X_{k,1}, \dots, X_{k,\nu-1})$ is past local information for the k -th data stream, and $\mathbf{P}^{(\nu)}$ and $\mathbf{E}^{(\nu)}$ denote the probability measure and expectation when the event occurs at time ν .

Mathematically, the problem of finding an efficient global monitoring scheme can then be formally formulated as finding a stopping time T such that the detection delay $\bar{\mathbf{E}}_{\delta_1, \dots, \delta_K}(T)$ is as small as possible for all possible combinations of $(\delta_1, \dots, \delta_K) \in \Delta$ in (24) subject to the false alarm constraint (2).

We are now ready to present the asymptotic optimality properties of our proposed schemes, $N_{hard}(a, b)$, $N_{soft}(a, b)$, $N_{order,r}(a)$, and $N_{comb,r}(a, b)$, under the standard asymptotic setting in which the number of data streams K is fix and the false alarm constraint γ goes to ∞ . Later we will briefly add some general remarks, including the properties when both the number K of data streams and the false alarm constraint γ go to ∞ in some appropriate rates.

The following theorem, whose proof is postponed to Section 5, derives the information bound on the detection delays of any globally monitoring schemes when Δ is defined in (24), as the false alarm constraint γ in (2) goes to ∞ .

Theorem 3.1. *Assume a scheme $T(\gamma)$ satisfies the false alarm constraint (2). Then for any given post-change hypothesis $(\delta_1, \dots, \delta_K) \in \Delta$, as γ goes to ∞ ,*

$$\bar{\mathbf{E}}_{\delta_1, \dots, \delta_K}(T(\gamma)) \geq (1 + o(1)) \frac{\log \gamma}{J(\delta_1, \dots, \delta_K)}, \quad (25)$$

where

$$J(\delta_1, \dots, \delta_K) = \sum_{k=1}^K I(g_k, f_k) I\{\delta_k < \infty\}, \quad (26)$$

and $I(g_k, f_k)$ is the KL information number defined in (12), and $I\{A\}$ is the indicator function of set A .

Next, when the local detection statistics $W_{k,n}$'s are the local CUSUM statistics in (14) and the local censoring parameters are defined by $b_k = \rho_k b$'s in (21)-(22) for some common constant $b \geq 0$, we establish

the asymptotic properties of our proposed schemes, $N_{hard}(a, b)$ in (16), $N_{order,r}(a)$ in (19), and $N_{comb,r}(a, b)$ in (18), as the global threshold a goes to ∞ , regardless of the false alarm constraint (2). The proof of the following theorem is presented in detail in Section 5.

Theorem 3.2. *As $a \rightarrow \infty$, let $b' = b'(a)$ be a constant such that both b' and $a - b'$ go to ∞ .*

(i) *The hard-thresholding scheme $N_{hard}(a, b)$ in (16) satisfies*

$$\mathbf{E}^{(\infty)}(N_{hard}(a, b)) \geq \frac{e^a}{1 + a + \frac{a^2}{2!} + \dots + \frac{a^{K-1}}{(K-1)!}}. \quad (27)$$

for any real number $b \geq 0$. Moreover, for any combination $(\delta_1, \dots, \delta_K) \in \Delta$ defined in (24), and for all $0 \leq b \leq b'$, we have

$$\overline{\mathbf{E}}_{\delta_1, \dots, \delta_K}(N_{hard}(a, b)) \leq \frac{a}{J(\delta_1, \dots, \delta_K)} + O(\sqrt{b}) + O(1) + O\left(\max_{\delta_k: \delta_k < \infty} (\delta_k)\right), \quad (28)$$

where the maximum is taken over all those δ_k 's that are finite, and $J(\delta_1, \dots, \delta_K)$ is defined in (26).

(ii) *The soft-thresholding scheme $N_{soft}(a, b)$ in (17) satisfies relation (27) for all $b \geq 0$. Moreover, it also satisfies relation (28) except that the a in the right-hand side of (28) is replaced by $a + b$.*

(iii) *For any integer $1 \leq r \leq K$, the order- r thresholding scheme $N_{order,r}(a)$ in (19) and the combined thresholding scheme $N_{comb,r}(a, b)$ in (18) with $b \geq 0$ also satisfy relation (27). In addition, for $0 \leq b \leq b'$, both schemes satisfy (28) whenever $\sum_{k=1}^K I\{\delta_k < \infty\} \leq r$, i.e., when the occurring event affects at most r sensors.*

Finally, when the local detection statistics $W_{k,n}$'s are the local CUSUM statistics in (14) and the local censoring parameters b_k 's are defined in (21)-(22) with $b \geq 0$ being the common constant, the asymptotic optimality properties of our proposed schemes can be summarized as follow.

Corollary 3.1. *For a given K and for any $b \geq 0$, with the choice of*

$$a = a_\gamma = \log \gamma + (K - 1 + o(1)) \log \log \gamma, \quad (29)$$

the hard-thresholding scheme $N_{hard}(a_\gamma, b)$ satisfies the false alarm constraint (2). Moreover, if $b' = b'_\gamma$ is chosen such that both b' and $a - b'$ go to ∞ as γ go to ∞ , then for all $0 \leq b \leq b'$,

$$\overline{\mathbf{E}}_{\delta_1, \dots, \delta_K}(N_{hard}(a, b)) \leq \frac{\log \gamma + (K - 1 + o(1)) \log \log \gamma}{J(\delta_1, \dots, \delta_K)} + O(\sqrt{b}) + O(1)$$

for all possible post-change hypothesis $(\delta_1, \dots, \delta_K) \in \Delta$ in (24). Therefore, for any given $b \geq 0$, the hard-thresholding schemes $N_{hard}(a, b)$ in (16) asymptotically minimize $\overline{\mathbf{E}}_{\delta_1, \dots, \delta_K}(N_{hard}(a, b))$ (up to the first-order) for each and every post-change hypothesis $(\delta_1, \dots, \delta_K) \in \Delta$ subject to the false alarm constraint (2), as γ in

(2) goes to ∞ . The conclusion also holds if $N_{hard}(a, b)$ is replaced by either the order-thresholding scheme $N_{order, r}$ in (19) or the combined thresholding scheme $N_{comb, r}(a, b)$ in (18) when the occurring event affects at most r data streams, i.e., when $(\delta_1, \dots, \delta_K) \in \Delta$ satisfies $\sum_{k=1}^K I\{\delta_k < \infty\} \leq r$.

Proof: This corollary follows at once from Theorems 3.1 and 3.2. In particular, the choice of a_γ in (29) follows from (27) and the fact that $1 + a + \frac{a^2}{2!} + \dots + \frac{a^{K-1}}{(K-1)!} \sim \frac{a^{K-1}}{(K-1)!}$ if K is fixed and a goes to ∞ . \square

It is worth pointing out several implications of our asymptotic results. First of all, from Corollary 3.1, it is interesting to note that the first-order term of the detection delays of the hard-thresholding scheme $N_{hard}(a_\gamma, b_\gamma)$ is $(\log \gamma)/J(\delta_1, \dots, \delta_K)$, the asymptotic lower bound in (25) in Theorem 3.1, but its second-order term contains both $O(\log \log \gamma)$ and $O(\sqrt{b})$. Hence, as the common constant b changes from 0 to $b'_\gamma = \log \gamma (\sim a_\gamma)$, the second-order term of the detection delays changes from $O(\log \log \gamma)$ to $O(\sqrt{\log \gamma})$. Hence, if we want to keep the second-order term of the detection delay to be as small as the order of $O(\log \log \gamma)$ for each and every possible post-change hypothesis (i.e., different combination of affected data streams), then the maximum choice of b should be $b = O((\log \log \gamma)^2) = O((\log a_\gamma)^2)$.

Second, recall that relations (10) and (11) provide heuristic choices of the global threshold a based on Chebyshev's inequality and the CLT approximation, respectively. For the purpose of better understanding these heuristic choices, below we will apply the spirit of these approximations to relation (27) when K is large. Note that the right-hand side of (27) is just $1/\mathbf{P}(U_K \geq a)$, where U_K denotes the sum of K iid exponential random variables with mean 1. To estimate the small value $\mathbf{P}(U_K \geq a)$ for large K , one way is to use the CLT that leads to $(U_K - K)/\sqrt{K} \sim N(0, 1)$. A choice of $a_\gamma \approx K + z_{1/\gamma}\sqrt{K}$ will yield $\mathbf{P}(U_K \geq a) = 1/\gamma$, and thus the right-hand side of (27) satisfies the global false alarm constraint in (2). This is consistent with the heuristic choice of a in (11).

The other way is to use Chebyshev's inequality and the theory of large deviations: for any constant $w = a/K > 1$, we have

$$\lim_{K \rightarrow \infty} -\frac{1}{K} \log \mathbf{P}(U_K \geq a) = \lim_{K \rightarrow \infty} -\frac{1}{K} \log \mathbf{P}\left(\frac{1}{K} U_K \geq w\right) = w - 1 - \log(w),$$

see, for example, Durrett [6, Ch. 1.9]. Hence, when the global false alarm constraint γ in (2) and the dimension K go to ∞ simultaneously in such a way that $(\log \gamma)/K = w - 1 - \log(w)$ is constant, a choice of the threshold $a = wK = (\log \gamma)w/(w - 1 - \log w)$ will lead the right-hand side of (27) satisfy the global false alarm constraint in (2). This is similar to the choice of a in (10).

We should mention that given the above heuristic choices of a , a comparison of relation (28) in Theorem 3.2 with the lower bound in Theorem 3.1 for fixed K and large γ suggests that our proposed schemes may no longer achieve the lower bound in Theorem 3.1, which may or may not provide a sharp lower bound on the detection delays as both the dimension K and the global false alarm constraint γ go to ∞ simultaneously in a suitable rate.

Third, let us further elaborate the communication rate between sensors and the fusion center in the context of monitoring K homogeneous sensors. As mentioned in (23), if we want at most $100\eta\%$ of K homogeneous sensors on average to transmit messages to the fusion center at any given time when no event occurs, we can choose the local censoring parameter $b_k = \log \eta^{-1}$, and thus the common censoring constant of our proposed schemes will be $b = K \log \eta^{-1}$ for a given $\eta \in (0, 1)$. Meanwhile, in our theorems and corollary, the asymptotic optimality properties of our proposed schemes hold under the condition that $(a - b)$ goes to ∞ . When the global threshold a satisfies $a = a_\gamma \approx K + \sqrt{K} z_{1/\gamma}$ as stated in the previous remark, then the condition of $a - b \rightarrow \infty$ is equivalent to $\log \eta^{-1} \leq 1$, i.e., $\eta \geq 1/e = 36.8\%$. In other words, when at least 36.8% of K homogeneous sensors can transmit messages at any given time when no event occurs, we can still develop efficient global monitoring schemes (e.g., $\{N_{hard}(a, b)\}$) that are asymptotically optimal to detect each and every possible combination of affected data streams. However, if $\eta < 36.8\%$, then it is unclear whether our proposed schemes can still effectively detect all different possible post-change hypotheses. Also see our numerical simulations in the previous section.

4 A Second Example: Normal Distribution with Unknown Post-Change Means

Suppose that we are monitoring K data streams $X_{k,n}$'s in (1). Initially, the data $X_{k,n}$'s are iid $N(0, 1)$. At some unknown time ν , an occurring event may change the distribution of the k -th local data stream to $N(\mu_k, 1)$ for some unknown $1 \leq k \leq K$. As in the previous section, we do not know which subset of local data streams are affected, but here we add a new challenge that we do not know the values of the post-change means μ_k 's when affected. We want to develop a system-wise online monitoring scheme that can detect the change as soon as possible, subject to the global false alarm constraint γ in (2).

Xie and Siegmund [38] investigates this problem under the assumption that the post-change mean $\mu_k > 0$ for all k . By assuming that the fraction p_0 of affected data stream is known, the main scheme they proposed is motivated from a semi-Bayesian approach and is defined by

$$T_{XS}(a, p_0) = \inf \left\{ n \geq 1 : \max_{0 \leq i < n} \sum_{k=1}^K \log(1 - p_0 + p_0 \exp \left[(U_{k,n,i}^+)^2 / 2 \right] \geq a \right\}. \quad (30)$$

where for all $1 \leq k \leq K, 0 \leq i < n$,

$$U_{k,n,i}^+ = \max \left(0, \frac{1}{\sqrt{n-i}} \sum_{j=i+1}^n X_{k,j} \right)$$

Some simplified versions have also been proposed to reduce the memory requirement to a large window of the most recent observations. However, all schemes in Xie and Siegmund [38] are not suitable in the context

of censoring sensor networks in Figure 1: besides being computationally expensive, the implementation of their schemes requires the fusion center to have full access to all data streams at each time step.

It has been an open problem to develop a scalable global monitoring scheme in the censoring sensor networks context that can detect both positive and negative local mean shifts for affected local data streams. Part of the reason is that for the K local data streams, there are 2^K potential different combinations of positive or negative local shifts, which is huge for a large K .

In this section, we illustrate how to tackle this open problem based upon our proposed SUM-shrinkage statistics in (5). The main challenge is to choose a suitable local detection statistic $W_{k,n}$ that can be easily computed and has the ability to detect both positive and negative local mean shifts. Once such local detection statistic $W_{k,n}$'s are defined, it is evident from the previous section that we can use any shrinkage transformation such as hard-thresholding, soft-thresholding, or order-thresholding to develop a global monitoring scheme. Below we use the soft-thresholding transformation as a demonstration. Our numerical simulation experience suggests that as a continuous function, the soft-thresholding transformation often yields smaller detection delays than the hard-thresholding transformation, and is computationally more efficient than the order-thresholding transformation.

To be more specific, in this section we will consider the soft-thresholding scheme

$$N_{soft}(a) = \inf \left\{ n \geq 1 : \sum_{k=1}^K \max(W_{k,n} - b_1, 0) \geq a \right\}, \quad (31)$$

where, for simplicity, all transformations h_k 's are chosen to the same soft-thresholding transformation $\max(u - b_1, 0)$ for some constant $b_1 > 0$. Our focus is how we can construct the local detection statistics $W_{k,n}$'s suitably.

The remainder of this section is as follows. Subsection 4.1 reviews the recursive register approach of Lorden and Pollak [18] for monitoring a single data stream, which is adapted to monitoring positive and negative mean shifts in Subsection 4.2. Subsection 4.3 provides a Bayesian interpretation of the soft-thresholding scheme as well as an efficient numerical algorithm of our proposed SUM-shrinkage scheme that only uses fixed $6K$ registers to store all past information and involves $O(K)$ computations at each given time step n . Numerical simulation results are summarized in subsection 4.4.

4.1 The Recursive Register Approach of Lorden and Pollak [18]

To abuse notation, in this subsection we suppress the subscript k of the k -th data stream, and consider the local monitoring problem with respect to the one-dimensional data stream $\{X_1, X_2, \dots\}$ whose distribution may change from $N(0, 1)$ to $N(\mu, 1)$ with unknown post-change mean μ at some unknown time ν . Lorden and Pollak [18] focuses on the case when the unknown post-change mean $\mu > 0$, and makes a technical assumption that $\mu \geq \rho$, where $\rho \geq 0$ is the smallest mean shift that is meaningful in practice, e.g. $\rho = 0.25$.

A high-level description of the recursive register approach of Lorden and Pollak [18] is as follows. Recall that the CUSUM statistics are defined in (14), and for one-dimensional normal distributed data, the CUSUM statistics have a simpler recursive form

$$W_n = \max(W_{n-1} + \mu X_n - \frac{1}{2}\mu^2, 0) \quad (32)$$

and $W_0 = 0$. When μ is unknown, we can continue to use this recursive formula to define a detection statistic if we replace the true unknown μ by its estimate from the past observed data. A key observation in Lorden and Pollak [18] is that at each given time step n , the CUSUM-type detection statistics can produce a candidate post-change time $\hat{\nu} \in \{0, 1, \dots, n-1\}$, and thus the observations $X_{\hat{\nu}}, X_{\hat{\nu}+1}, \dots, X_{n-1}$ can be used to estimate the post-change mean μ in (32). Specifically, at any given time step n , define $\hat{\nu}$ as the largest $0 \leq i \leq n-1$ such that $W_i = 0$, and denote by T_n and S_n the total number and the summation of observations X_i 's between the candidate post-change time $\hat{\nu}$ and time step $n-1$. That is,

$$T_n = n - \hat{\nu} \quad \text{and} \quad S_n = \sum_{i=\hat{\nu}}^{n-1} X_i. \quad (33)$$

By the method of moments estimator or maximum likelihood estimator method, the post-change mean μ can be estimated by S_n/T_n at time step n . If we treat the pre-specified nonnegative constants s and t as a prior, then a Bayes-type estimate of μ is $\hat{\mu}_n = (s + S_n)/(t + T_n)$, which includes S_n/T_n as a special case when $s = t = 0$. After taking into account that ρ is the smallest post-change mean we are interested in, one can estimate μ at time step n by

$$\hat{\mu}_n = \max\left(\rho, \frac{s + S_n}{t + T_n}\right). \quad (34)$$

From the algorithm viewpoint, the recursive register approach of Lorden and Pollak [18] can be recursively implemented as follows. Let $S_0 = T_0 = W_0 = X_0 = 0$, and $\hat{\mu}_1 = \rho$. For all $n \geq 1$,

$$W_n = \max\left(W_{n-1} + \hat{\mu}_n X_n - \frac{1}{2}(\hat{\mu}_n)^2, 0\right), \quad (35)$$

where $\hat{\mu}_n$ is defined in (34) and

$$\begin{pmatrix} S_n \\ T_n \end{pmatrix} = \begin{cases} \begin{pmatrix} S_{n-1} + X_{n-1} \\ T_{n-1} + 1 \end{pmatrix} & \text{if } W_{n-1} > 0 \\ \begin{pmatrix} 0 \\ 0 \end{pmatrix} & \text{if } W_{n-1} = 0. \end{cases} \quad (36)$$

In other words, the local detection statistics W_n 's can be computed recursively as the part of three-dimensional vectors (S_n, T_n, W_n) , or four-dimensional vectors $(S_n, T_n, \hat{\mu}_n, W_n)$. It is important to note that $(S_n, T_n, \hat{\mu}_n)$ only uses the observations up to time $n-1$ for the purpose of estimating the post-change mean

μ , so that the data X_n is reserved for the local detection statistics W_n for the purpose of detecting changes. It was shown that the detection scheme based on the detection statistic W_n in (35) is asymptotically optimal whenever the true post-change mean $\mu \geq \rho > 0$, see Theorems 3.1-3.3 of Lorden and Pollak [18].

4.2 Our Proposed Local Detection Statistics $W_{k,n}$'s

Since we are interested in detecting both positive and negative local mean shifts for affected data streams, we propose to extend the detection statistic W_n in (35) of Lorden and Pollak [18] from one-sided to two-sided. Observe that detecting negative local mean shift of $X_{k,n}$'s is equivalent to detecting positive local mean shift of $-X_{k,n}$'s, we propose the following two-sided local detection statistic for each local data stream at time n :

$$W_{k,n} = \max(W_{k,n}^{(1)}, W_{k,n}^{(2)}), \quad (37)$$

where $W_{k,n}^{(1)}$ and $W_{k,n}^{(2)}$ are the local detection statistics of Lorden and Pollak [18] for detecting positive and negative mean shifts, respectively. Specifically,

$$\begin{aligned} W_{k,n}^{(1)} &= \max \left(W_{k,n-1}^{(1)} + \hat{\mu}_{k,n}^{(1)} X_{k,n} - \frac{1}{2} (\hat{\mu}_{k,n}^{(1)})^2, 0 \right), \\ W_{k,n}^{(2)} &= \max \left(W_{k,n-1}^{(2)} + \hat{\mu}_{k,n}^{(2)} X_{k,n} - \frac{1}{2} (\hat{\mu}_{k,n}^{(2)})^2, 0 \right), \end{aligned} \quad (38)$$

where

$$\hat{\mu}_{k,n}^{(1)} = \max \left(\rho, \frac{s + S_{k,n}^{(1)}}{t + T_{k,n}^{(1)}} \right) > 0, \quad \hat{\mu}_{k,n}^{(2)} = \min \left(-\rho, \frac{-s + S_{k,n}^{(2)}}{t + T_{k,n}^{(2)}} \right) < 0, \quad (39)$$

and for $j = 1, 2$ and for any k , the sequences $(S_{k,n}^{(j)}, T_{k,n}^{(j)})$ are defined recursively

$$\begin{pmatrix} S_{k,n}^{(j)} \\ T_{k,n}^{(j)} \end{pmatrix} = \begin{cases} \begin{pmatrix} S_{k,n-1}^{(j)} + X_{k,n-1} \\ T_{k,n-1}^{(j)} + 1 \end{pmatrix} & \text{if } W_{k,n-1}^{(j)} > 0 \\ \begin{pmatrix} 0 \\ 0 \end{pmatrix} & \text{if } W_{k,n-1}^{(j)} = 0 \end{cases} \quad (40)$$

Note that $\hat{\mu}_{k,n}^{(1)}$ and $\hat{\mu}_{k,n}^{(2)}$ in (39) are the estimates of the post-change mean when restricted to the positive and negative values, respectively, under the assumption that $|\mu| \geq \rho$. Clearly, $W_{k,n}^{(1)}$ is designed to detect positive local mean shift, whereas $W_{k,n}^{(2)}$ is to detect negative local mean shifts. Also the two-sided local detection statistic $W_{k,n}$ in (37) is always nonnegative for any k at any time step n , and it will become large when there is a local mean shift no matter whether such mean shift is positive or negative.

4.3 Interpretation and Overview of The Soft-Thresholding Scheme

With the local detection statistics $W_{k,n}$'s in (37), the soft-thresholding scheme in (31) can be used to monitor K data streams with possible local positive or negative mean shifts. It is natural to ask why the soft-

thresholding scheme in (31) works? Besides the motivation in the previous section, below we also provide a semi-Bayesian interpretation.

At a given time n , let Z_k be the indicator whether the distribution of the k -th local data stream changes for $k = 1, \dots, K$. Assume that each local data stream has a prior probability π getting affected by the event, and assume that Z_1, \dots, Z_K are iid with probability mass function $\mathbf{P}(Z_k = 1) = \pi = 1 - \mathbf{P}(Z_k = 0)$. Treat Z_k 's as the hidden states, and recall that $W_{k,n}$ represents the evidence of possible change (in logarithm scale) and is applicable only when $Z_k = 1$ (since $Z_k = 0$ implies that there is no change at the k -th data stream). Then when testing $H_0 : Z_1 = \dots = Z_K = 0$ (no change), the log-likelihood ratio (LLR) statistic of the hidden state Z_k 's and the observed data $X_{k,n}$'s is

$$\begin{aligned} LLR(n) &= \sum_{k=1}^K \{Z_k(\log \pi + W_{k,n}) + (1 - Z_k) \log(1 - \pi)\} - \sum_{k=1}^K \log(1 - \pi) \\ &= \sum_{k=1}^K Z_k \{W_{k,n} - \log((1 - \pi)/\pi)\} \end{aligned}$$

Since the Z_k 's are unobservable, it is natural to maximize $LLR(n)$ over $Z_1, \dots, Z_K \in \{0, 1\}$. Hence, the maximum likelihood estimator of the Z_k 's is that

$$\hat{Z}_k = \begin{cases} 1, & \text{if } W_{k,n} \geq \log((1 - \pi)/\pi) \\ 0, & \text{otherwise} \end{cases}, \quad \text{for } k = 1, \dots, K,$$

and the generalized log-likelihood ratio becomes

$$\max_{Z_k's} LLR(n) = \sum_{k=1}^K \max\{W_{k,n} - \log((1 - \pi)/\pi), 0\},$$

which is exactly the form of the soft-thresholding scheme $N_{soft}(a)$ in (31) with $b_1 = \log((1 - \pi)/\pi)$.

The above discussion indicates that if we have a prior knowledge that a fraction of π sensors will be affected by the event, the censoring parameter b_1 can be chosen as $b_1 = \log((1 - \pi)/\pi)$. Meanwhile, as in the previous section, if we want η proportion of data streams to send information to the fusion center when no change occurs, one rule of thumb is to choose $b_1 = \log \eta^{-1}$. It is interesting to see that these two values of b_1 are very close when $\eta = \pi$ is small. In practice, the true value of π is often unknown, but one may have a preferred η value. Hence, in our simulations below we will choose the censoring parameter $b_1 = \log \eta^{-1}$.

Note that the proposed soft-thresholding scheme $N_{soft}(a)$ in (31) can be easily implemented in the censoring sensor network context by parallel computing the K local detection statistics $W_{k,n}$'s recursively through (37)-(40) at the local sensor levels. To be more specific, we can use the following $6K$ registers to adaptively store all past information at each time step after observing new data: $(S_k^{(j)}, T_k^{(j)}, W_k^{(j)})$ for $j = 1, 2$ and $k = 1, 2, \dots, K$. At any given time step n , we can first update the $4K$ registers in $(S_k^{(j)}, T_k^{(j)})$ using the past data and compute the $2K$ estimates $\hat{\mu}_k^{(j)}$ of the post-change means μ_k 's. Then after we

observe new observations, $(X_{1,n}, \dots, X_{K,n})$, we only need to update the $2K$ registers $W_k^{(j)}$'s and compute the values of K local detection statistics W_k 's, which allows us to easily compute the global monitoring statistic G . Including the $3K$ intermediate variables $(\hat{\mu}_k^{(j)}, W_k)$ and the global monitoring statistic G , the proposed scheme only needs $9K + 1$ registers to adaptively store all relevant information and involves $O(K)$ computations at any given time step n . Moreover, our proposed scheme can be implemented in the context of censoring sensor networks in the previous section where most computations are done at the remote sensors and the communication cost and the computational burden at the fusion center are marginal. Hence, our proposed scheme is scalable and can be easily implemented to online monitor large-scale data streams over a long time period.

An overview of our proposed N_{soft} scheme is illustrated in the following algorithm:

Algorithm: Implementation of N_{soft} in (31)

Initial parameters: ρ , s , t , and b_1 for $k = 1, \dots, K$.

Set: A terminal threshold a .

Algorithm:

initialize $n = 0$, and set all initial observations $X_k = 0$ and all $8K$ initial registers $S_k^{(j)} = T_k^{(j)} = \mu_k^{(j)} = W_k^{(j)} = 0$, for $k = 1, \dots, K$ and $j = 1, 2$.

While the scheme N_{soft} has not raised an alarm

- do**
1. Update $4K$ registers $(S_k^{(j)}, T_k^{(j)})$ via (40).
 2. Compute the $2K$ intermediate variables $\hat{\mu}_k^{(j)}$ from (39) which are the estimates of the post-change means.
 3. Input new observations from all K data streams, denoted by (X_1, \dots, X_K) .
 4. For $k = 1, \dots, p$, recompute the local monitoring statistics $W_k^{(j)}$'s in (38) and W_k in (37).
 5. Compute the global monitoring statistics

$$G = \sum_{k=1}^K \max(W_k - b_1, 0)$$

if $G \geq a$ **terminate:** Raising an alarm at time n and declaring that a change has occurred;

end the while loop

4.4 Simulation Results

In this section, we report the numerical simulation results of the soft-thresholding scheme $N_{soft}(a)$ in (31). For the purpose of comparison, we follow Xie and Siegmund [38] to assume that there are $K = 100$ indepen-

dent normal data streams. For each $k = 1, \dots, K$, the data $X_{k,n}$'s of the k -th data stream are iid $N(0, 1)$ before the change, but are iid $N(1, 1)$ after the k -th data stream is affected by the occurring event.

In our simulations, we consider six schemes: two of them are the Xie and Siegmund schemes $T_{XS}(a, p_0)$ in (30) with $p_0 = 1$ and 0.1 ; and the remaining four schemes are our proposed soft-thresholding schemes $N_{soft}(a)$ in (31) with four different thresholding parameters: $b_1 = 0, 0.5, \log(10), \log(100)$. As in the previous section, the three non-zero b_1 values imply that on average at most $\exp(-b_1) \approx 60.1\%, 10\%$ and 1% out of 100 local data streams produce significant $W_{k,n}$'s values to the global monitoring statistic G_n when there are no changes. When computing the local detection statistics $W_{k,n}$'s in (37), we set $\rho = 0.25, t = 4$ and $s = 1$ as in Lorden and Pollak [18].

For each of these six schemes $T(a)$, we first numerically search the threshold a to satisfy the global false alarm constraint γ in (2). Two different values of γ are considered. One is $\gamma = 5000$, so that we can compare with those results from Xie and Siegmund [38]. The other is $\gamma = 5 \times 10^4$ to see the effect of false alarm constraint γ on the detection delays of our proposed schemes. Note that we are unable to numerically find the global threshold a of the Xie and Siegmund scheme for the case of $\gamma = 5 \times 10^4$ in a reasonable time, and thus we will only report the performance of our proposed schemes. Next, for the detection delays of $T(a)$, we consider various post-change hypotheses, and for each post-change hypothesis, we simulate the $\mathbf{E}(T(a))$ when the event occurs at time $\nu = 1$, and use this as an estimate of the detection delay $\mathbf{D}(T(a))$. All simulated values are based on 2500 Monte Carlo runs.

Table 2 summarizes the detection delays in the scenario when the change is instantaneous if a local data stream is affected. For the Xie and Siegmund scheme $T_{XS}(a, p_0)$ in (30), our simulated detection delay results are slightly different from their reported results in their paper, possibly because our simulation is based on 2500 runs instead of 500 runs in their paper. Note that the Xie and Siegmund schemes $T_{XS}(a, p_0)$ in (30) involve expensive computations, and require the fusion center to have full access to all raw data. Thus it is not surprising that their schemes have smaller detection delays than our proposed soft-thresholding schemes. However, we want to emphasize that the Xie and Siegmund schemes are not scalable and cannot be implemented in the context of distributed monitoring in censoring sensor networks. Meanwhile, our proposed schemes are suitable to the censoring sensor network contexts, as they can be easily implemented by parallel computing in a recursive manner at the local sensors level and the computational costs between the local sensors and the fusion center will be marginal.

A more reasonable comparison is to compare the results in Table 2 with those in Table 1 which were conducted under the assumption that the post-change mean of each affected local data stream is $\mu = 1$. When at least 5 local data streams are affected, the detection delays of $N_{soft}(a = 24.01, b_1 = \log(10))$ in Table 2 are only $2 \sim 5$ larger than those of $N_{soft}(a = 21.56, b_k = 2.3026)$ in Table 1. Since the schemes in Table 2 are able to detect both positive or negative mean shifts, one may be willing to pay the price of

Table 2: A comparison of detection delays when the change is instantaneous and the post-change mean $\mu_k = 1$ if affected. The smallest and largest standard errors of the schemes are also reported under each post-change hypothesis based on 2500 repetitions in Monte Carlo simulations.

γ		# local data streams affected								
		1	3	5	8	10	20	30	50	100
	Smallest standard error	0.19	0.08	0.06	0.04	0.03	0.02	0.01	0.01	0.00
	Largest standard error	0.40	0.14	0.08	0.05	0.04	0.03	0.02	0.02	0.01
5000	Xie and Siegmund's schemes $T_{XS}(a, p_0)$ in (30)									
	$T_{XS}(a = 53.5, p_0 = 1)$	52.4	18.3	11.1	7.1	5.7	2.9	2.0	1.2	1.0
	$T_{XS}(a = 19.5, p_0 = 0.1)$	31.1	13.4	9.2	6.7	5.7	3.5	2.5	1.8	1.0
	Soft-thresholding Schemes $N_{soft}(a)$ in (31)									
	$N_{soft}(a = 127.86, b_1 = 0)$	75.0	35.4	25.2	18.5	16.0	10.3	8.1	6.1	4.1
	$N_{soft}(a = 84.91, b_1 = 0.5)$	72.1	33.9	24.1	17.7	15.3	10.0	7.9	6.0	4.2
	$N_{soft}(a = 24.01, b_1 = \log(10))$	45.8	22.0	16.4	12.8	11.5	8.5	7.3	6.1	5.0
	$N_{soft}(a = 7.88, b_1 = \log(100))$	29.0	17.2	14.2	12.0	11.2	9.2	8.3	7.3	6.4
5×10^4	Soft-thresholding Schemes $N_{soft}(a)$ in (31)									
	$N_{soft}(a = 136.07, b_1 = 0)$	89.0	39.9	27.9	20.2	17.4	11.1	8.7	6.5	4.4
	$N_{soft}(a = 92.79, b_1 = 0.5)$	85.7	38.2	26.8	19.4	16.7	10.7	8.4	6.3	4.4
	$N_{soft}(a = 29.05, b_1 = \log(10))$	55.1	25.3	18.4	14.1	12.6	9.1	7.8	6.5	5.2
	$N_{soft}(a = 11.11, b_1 = \log(100))$	35.5	19.7	16.0	13.4	12.4	10.0	8.9	7.9	6.8

slightly larger detection delays at the given post-change mean $\mu = 1$ so as to effectively detect other local mean shifts, especially the negative shifts. In addition, it is interesting to see from Table 2 that as the false alarm constraint γ increases from 5000 to 5×10^4 , the global threshold a of our proposed soft-thresholding schemes $N_{soft}(a, b_1)$ increases moderately for any given censoring parameters b_1 , but the detection delays of our proposed soft-thresholding schemes increase only marginally when at least 5 data streams are affected.

All simulations were done on a Windows 8 Laptop with Intel i7-4700MQ CPU 2.40GHz using MATLAB R2013b. For each of these schemes $T(a)$ (i.e., each row of Table 2), the most time consuming part was to search for the global threshold a so that $\mathbf{E}^{(\infty)}(T(a)) \approx \gamma$. When $\gamma = 5000$, it took about 8 minutes to find such a from a *range of values* for our proposed schemes based on 2500 Monte Carlo runs (the time is shorter if our initial guess range of a is closer). Meanwhile, for the Xie and Siegmund scheme, for a *given* global threshold a around 53.5 which was provided in their paper, it took about one and a half hour on average to finish one Monte Carlo simulation run in our laptop. If we did not know $a \approx 53.5$ and wanted

to try 10 different values of a 's by bisection method based on 2500 Monte Carlo runs for each a , it would have taken about $10 \times 1.5 \times 2500 = 37500$ computer hours for the case of $\gamma = 5000$. When $\gamma = 5 \times 10^4$, it took us about one hour to find the global threshold a for our proposed schemes, but we are unable to numerically implement the Xie and Siegmund schemes since their computational time will be in days for each Monte carlo run. Once the global threshold a is found, it is straightforward to simulate the detection delays in Table 2. When $\gamma = 5000$, our proposed schemes are at least 10 times faster than the Xie and Siegmund schemes. For instance, when exactly one data stream is affected, it took 4.94 seconds to simulate the detection delay of our proposed schemes, whereas it took 41.02 seconds to simulate those of the Xie and Siegmund schemes. Hence, as compared to the Xie and Siegmund schemes, the computational advantage of our proposed schemes is evident.

5 Proof of Theorems 3.1 and 3.2

This section is devoted to prove Theorems 3.1 and 3.2.

Proof of Theorem 3.1. Intuitively, only those affected sensors provide information to detect the occurring events, and the quickest possible way to detect the occurring event is when the event affects the sensors instantaneously. More rigorously, if we define

$$\delta_k^* = \begin{cases} 0, & \text{if } \delta_k \text{ is finite} \\ \infty, & \text{if } \delta_k = \infty \end{cases}, \quad (41)$$

then for any given scheme $T(\gamma)$,

$$\overline{\mathbf{E}}_{\delta_1, \dots, \delta_K}(T(\gamma)) \geq \inf_{\tau} \overline{\mathbf{E}}_{\delta_1^*, \dots, \delta_K^*}(\tau),$$

where the infimum is taken over all possible schemes τ satisfying the false alarm constraint γ in (2). An alternative and possible better viewpoint is based on a time-shifting argument in which one imagines that at time n one observes the observations $X_{k, n+\delta_k}$ (instead of $X_{k, n}$) when δ_k is finite, and then applies $T(\gamma)$ to the new aligned observations.

Without loss of generality, assume that the first m data streams are affected abruptly and simultaneously by the event at unknown time ν , and other data streams are unaffected. That is, m out of K data streams are affected by the event, and $\delta_i^* = 0$ for $1 \leq i \leq m$, and $= \infty$ for $m+1 \leq i \leq K$. By (26), we have

$$J(\delta_1, \dots, \delta_K) = J(\delta_1^*, \dots, \delta_K^*) = \sum_{i=1}^m I(g_i, f_i).$$

In this case, we face the sequential change detection problem when the distribution of $(X_{1, n}, \dots, X_{K, n})$ changes from $(f_1, \dots, f_m, f_{m+1}, \dots, f_K)$ to $(g_1, \dots, g_m, f_{m+1}, \dots, f_K)$. It is well-known (Lorden [17]) that

$$\inf_{\tau} \overline{\mathbf{E}}_{\delta_1^*, \dots, \delta_K^*}(\tau) \geq (1 + o(1)) \frac{\log \gamma}{\sum_{i=1}^m I(g_i, f_i)}.$$

subject to the false alarm constraint γ in (2) as $\gamma \rightarrow \infty$. Combining the above results yields relation (25), completing the proof of Theorem 3.1. \square

Proof of Theorem 3.2. Let us first focus part (i) on the properties of the hard-thresholding scheme $N_{hard}(a, b)$ in (16) with $b \geq 0$ being the common constant for b_k 's in (21)-(22).

To prove (27), note that $N_{hard}(a, b)$ in (16) is increasing as a function of $b \geq 0$, and when $b = 0$, $N_{hard}(a, b = 0)$ reduces to the ‘‘SUM’’ scheme $T_{sum}(a)$ in (4). Hence, for any $b \geq 0$, $N_{hard}(a, b) \geq T_{sum}(a)$ and of course, $\mathbf{E}^{(\infty)}(N_{hard}(a, b)) \geq \mathbf{E}^{(\infty)}(T_{sum}(a))$. By Theorem 1 of Mei [20], the ‘‘SUM’’ scheme $T_{sum}(a)$ satisfies relation (27), and so are the hard-thresholding schemes $N_{hard}(a, b)$ for all $b \geq 0$.

To prove relation (28), it is clear that the worst-case detection delay of $N_{hard}(a, b)$ occurs at the change-point $\nu = 1$, and thus it suffices to show that $\mathbf{E}_{\delta_1, \dots, \delta_K}^{(\nu=1)}(N_{hard}(a, b))$ satisfies (28). Without loss of generality, we assume that only the first m data streams are affected and no other data streams are affected. To simplify our notation below, denote $\delta_{\max} = \max_{1 \leq i \leq m} \delta_i$. It suffices to show that

$$\mathbf{E}_{\delta_1, \dots, \delta_K}^{(\nu=1)}(N_{hard}(a, b)) \leq \frac{a}{\sum_{k=1}^m I(g_k, f_k)} + O(\sqrt{b}) + O(1) + \delta_{\max}, \quad (42)$$

for all $0 \leq b \leq b'$ when b' and $(a - b')$ go to ∞ .

The essential idea in the proof of (42) is to compare $N_{hard}(a, b)$ with new stopping times that are only based on those affected m data streams. Define a stopping time that is in the form of the one-sided sequential probability ratio test (SPRT):

$$\tau(a, b) = \text{first } n \text{ such that } \sum_{i=1}^n \sum_{k=1}^m \log \frac{g_k(X_{k,i})}{f_k(X_{k,i})} \geq a \text{ and } \sum_{i=1}^n \log \frac{g_k(X_{k,i})}{f_k(X_{k,i})} \geq \rho_k b \text{ for all } 1 \leq k \leq m, \quad (43)$$

where the weights ρ_k 's are defined in (22), and let $\hat{\tau}_\delta(a, b)$ be the new stopping time that applies $\tau(a, b)$ to the new observations after time δ_{\max} .

Now whenever $\hat{\tau}_\delta(a, b)$ stops at time $n_0 + \delta_{\max}$, we know that $\tau(a, b)$ stops after applying it to n_0 observations $(X_{k, \delta_{\max}+1}, \dots, X_{k, \delta_{\max}+n_0})$ for each k . By the definition of the local CUSUM statistics in (14), we have

$$W_{k, n_0 + \delta_{\max}} \geq \sum_{i=\delta_{\max}+1}^{\delta_{\max}+n_0} \log \frac{g_k(X_{k,i})}{f_k(X_{k,i})} \geq \rho_k b$$

for all $1 \leq k \leq m$. Hence,

$$\sum_{k=1}^K W_{k, n_0 + \delta_{\max}} \mathbf{1}\{W_{k, n_0 + \delta_{\max}} \geq \rho_k b\} \geq \sum_{k=1}^m \sum_{i=\delta_{\max}+1}^{\delta_{\max}+n_0} \log \frac{g_k(X_{k,i})}{f_k(X_{k,i})} \geq a,$$

where the last relation is from the definition of $\tau(a, b)$. This implies that the scheme $N_{hard}(a, b)$ must stop at time $n_0 + \delta_{\max}$, and possibly earlier. Thus

$$\mathbf{E}_{\delta_1, \dots, \delta_K}^{(\nu=1)}(N_{hard}(a, b)) \leq \mathbf{E}_{\delta_1, \dots, \delta_K}^{(\nu=1)}(\hat{\tau}_\delta(a, b)) = \delta_{\max} + \mathbf{E}_{\delta_1^*, \dots, \delta_K^*}^{(\nu=1)}(\tau(a, b)),$$

where δ_k^* is the binary version of δ_k 's defined in (41). To simplify the notation, denote by $\mathbf{E}^{(1)}$ the expectation when the change occurs at time $\nu = 1$ and the event affects the first m data streams immediately but does not affect the other remaining $K - m$ data streams. So it suffices to show that the stopping time $\tau(a, b)$ in (43) satisfies

$$\mathbf{E}^{(1)}(\tau(a, b)) \leq \frac{a}{\sum_{k=1}^m I(g_k, f_k)} + O(\sqrt{b}) + O(1). \quad (44)$$

To prove (44), for $1 \leq k \leq m$, let

$$\begin{aligned} M_k &= \inf \left\{ n \geq 1 : \sum_{i=1}^n \log \frac{g_k(X_{k,i})}{f_k(X_{k,i})} \geq \rho_k b \right\}, \\ \tau_k(M_k) &= \sup \left\{ n \geq 1 : \sum_{i=M_k+1}^{M_k+n} \log \frac{g_k(X_{k,i})}{f_k(X_{k,i})} \leq 0 \right\} \\ \hat{M} &= \max_{1 \leq k \leq m} (M_k + \tau_k(M_k) + 1) \\ t(\hat{M}) &= \inf \left\{ n \geq 1 : \sum_{i=\hat{M}+1}^{\hat{M}+n} \left(\sum_{k=1}^m \log \frac{g_k(X_{k,i})}{f_k(X_{k,i})} \right) \geq a - \left(\sum_{k=1}^m \rho_k \right) b \right\}. \end{aligned}$$

In the definition of $t(\hat{M})$, the assumption of $0 \leq b \leq b'$ is used to make sure that the threshold

$$a - \left(\sum_{k=1}^m \rho_k \right) b \geq a - b \geq a - b'$$

goes to ∞ as $a \rightarrow \infty$, since $\sum_{k=1}^m \rho_k \leq \sum_{k=1}^K \rho_k = 1$ and $(a - b')$ is assumed to go to ∞ . Combining these definitions with those of $\tau(a, b)$ in (43) yields that

$$\begin{aligned} \tau(a, b) &\leq \hat{M} + t(\hat{M}) = \max_{1 \leq k \leq m} (M_k + \tau_k(M_k) + 1) + t(\hat{M}) \\ &\leq \sum_{k=1}^m \tau_k(M_k) + 1 + t(\hat{M}) + \max_{1 \leq k \leq m} M_k. \end{aligned}$$

Hence, relation (44) holds if we can establish the following three relations:

$$\mathbf{E}^{(1)}(\tau_k(M_k)) = O(1) \quad \text{for all } 1 \leq k \leq m; \quad (45)$$

$$\mathbf{E}^{(1)}(t(\hat{M})) \leq \frac{a}{\sum_{k=1}^m I(g_k, f_k)} - \frac{b}{\sum_{k=1}^K I(g_k, f_k)} + O(1); \quad (46)$$

$$\mathbf{E}^{(1)}\left(\max_{1 \leq k \leq m} M_k\right) \leq \frac{b}{\sum_{k=1}^K I(g_k, f_k)} + O(\sqrt{b}) + O(1). \quad (47)$$

Relation (45) is well-known in renewal theory, e.g., Theorem D in Kiefer and Sacks [11], since $\log(g_k(X)/f_k(X))$ has positive mean and finite variance under $\mathbf{E}^{(1)}$ by our assumptions in (12) and (13).

For relation (46), by the definition of ρ_k in (22), we have

$$\frac{\sum_{k=1}^m \rho_k}{\sum_{k=1}^m I(g_k, f_k)} = \frac{1}{\sum_{k=1}^K I(g_k, f_k)}.$$

Since $t(\hat{M})$ is the stopping time when a random walk exceed the bound $a - (\sum_{k=1}^m \rho_k)b$, the application of standard renewal theory yields that

$$\begin{aligned}\mathbf{E}^{(1)}(t(\hat{M})) &= \frac{a - (\sum_{k=1}^m \rho_k)b}{\sum_{k=1}^m I(g_k, f_k)} + O(1) \\ &= \frac{a}{\sum_{k=1}^m I(g_k, f_k)} - \frac{b}{\sum_{k=1}^K I(g_k, f_k)} + O(1),\end{aligned}$$

as the threshold $a - (\sum_{k=1}^m \rho_k)b$ goes to ∞ , see, for example, Siegmund [32, Ch. VIII]. Thus relation (46) holds.

The proof of relation (47) is a little more complicated, but it can be done along the same line as that in Mei [19]. The key fact is that the choice of $b_k = \rho_k b$'s in (21)-(22) makes sure that the stopping times M_k 's have roughly the same mean under $\mathbf{P}^{(1)}$. Specifically, by renewal theory and the assumptions of (f_k, g_k) in (12) and (13), under $\mathbf{P}^{(1)}$,

$$\mathbf{E}^{(1)}(M_k) = \frac{\rho_k b}{I(g_k, f_k)} + O(1) = \frac{b}{\sum_{k=1}^K I(g_k, f_k)} + O(1)$$

and $\text{Var}^{(1)}(M_k) = O(b)$, as $b \rightarrow \infty$, see Siegmund [32, p. 171]. Thus

$$\begin{aligned}\left(\mathbf{E}^{(1)}\left|M_k - \frac{b}{\sum_{k=1}^K I(g_k, f_k)}\right|\right)^2 &\leq \mathbf{E}^{(1)}\left(M_k - \frac{b}{\sum_{k=1}^K I(g_k, f_k)}\right)^2 \\ &= \text{Var}^{(1)}(M_k) + \left(\mathbf{E}^{(1)}M_k - \frac{b}{\sum_{k=1}^K I(g_k, f_k)}\right)^2 \\ &= O(b)\end{aligned}$$

as $b \rightarrow \infty$. Hence, for each $k = 1, \dots, K$, there exist two constants $C_{1k} > 0$ and $C_{2k} > 0$ so that for all $b \geq 0$,

$$\left|M_k - \frac{b}{\sum_{k=1}^K I(g_k, f_k)}\right| \leq \max(C_{1k}, C_{2k}\sqrt{b}).$$

Therefore,

$$\begin{aligned}\mathbf{E}^{(1)}\left(\max_{1 \leq k \leq m} M_k\right) &= \frac{b}{\sum_{k=1}^K I(g_k, f_k)} + \mathbf{E}^{(1)}\max_{1 \leq k \leq m} \left(M_k - \frac{b}{\sum_{k=1}^K I(g_k, f_k)}\right) \\ &\leq \frac{b}{\sum_{k=1}^K I(g_k, f_k)} + \sum_{k=1}^m \mathbf{E}^{(1)}\left|M_k - \frac{b}{\sum_{k=1}^K I(g_k, f_k)}\right| \\ &\leq \frac{b}{\sum_{k=1}^K I(g_k, f_k)} + \sum_{k=1}^m \max(C_{1k}, C_{2k}\sqrt{b}) \\ &\leq \frac{b}{\sum_{k=1}^K I(g_k, f_k)} + C(\sqrt{b} + 1),\end{aligned}$$

where the constant $C = \sum_{k=1}^K \max(C_{1k}, C_{2k})$ does not depend on b . This proves relation (47). Therefore, relations (45)-(47) hold, and thus relation (28) holds for the hard-thresholding scheme $N_{hard}(a, b)$ in (16). The proof for the soft-thresholding scheme $N_{soft}(a, b)$ in (17) is identical and thus omitted.

Now let us provide a sketch of the proof for part (iii) of Theorem 3.2 on the order-thresholding scheme $N_{order,r}(a)$ in (19) and the combined thresholding scheme $N_{comb,r}(a, b)$ in (18). Since $N_{order,r}(a)$ is a special case of $N_{comb,r}(a, b)$ with $b = 0$, it suffices to prove the theorem for $N_{comb,r}(a, b)$ in (18) with $b \geq 0$. Clearly relation (27) also holds for $N_{comb,r}(a, b)$ for any $b \geq 0$, because the “SUM” scheme $T_{sum}(a)$ again provides the lower bound for $N_{comb,r}(a, b)$.

It remains to show that relation (28) holds for $N_{comb,r}(a, b)$ with $b \geq 0$ in the scenario when the occurring event affects at most r data streams, i.e., when $\sum_{k=1}^K I\{\delta_k < \infty\} \leq r$. Without loss of generality, assume that the affected data streams are just the first m data streams with $m \leq r$. Recall that $U_{k,n} = W_{k,n} I\{W_{k,n} \geq \rho_k b\}$, and we order the $U_{k,n}$ ’s as $U_{(1),n} \geq \dots \geq U_{(K),n}$, and $N_{comb,r}(a, b)$ stops if $\sum_{k=1}^r U_{(k),n} \geq a$. Note that if $m \leq r$,

$$\sum_{k=1}^r U_{(k),n} \geq \sum_{k=1}^r U_{k,n} \geq \sum_{k=1}^m U_{k,n},$$

since $U_{k,n} \geq 0$. Thus, if at some time n_0 we have $W_{k,n_0} \geq \rho_k b$ and $\sum_{k=1}^m W_{k,n_0} \geq a$ for $1 \leq k \leq m$ (i.e., for the first m data streams), then $N_{comb,r}(a, b)$ will also stop at time n_0 and possibly earlier. Hence, whenever $m \leq r$, the stopping time $\tau(a, b)$ in (43) also provides an upper bound on the detection delay of $N_{comb,r}(a, b)$. Thus the proposed combined thresholding scheme $N_{comb,r}(a, b)$ in (18) satisfies relation (28) whenever the occurring event affects at most r data streams. This completes the proof of the theorem. \square

References

- [1] APPADWEDULA, S., VEERAVALLI, V. V., and JONES, D. L. (2005). Energy-efficient detection in sensor networks. *IEEE J. Sel. Areas Commun.*, 23, 693–702.
- [2] BASSEVILLE, M. and NIKIFOROV, I. V. (1993). *Detection of Abrupt Changes: Theory and Applications*. Englewood Cliffs, Prentice-Hall. MR1210954
- [3] BREIMAN, L. (2001). Statistical modeling: the two cultures. *Statistical Sciences*, 16, 199–231. MR1874152
- [4] CANDÈS, E. J. (2006). Modern statistical estimation via oracle inequalities. *Acta Numerica*, 15, 257–325. MR2269743
- [5] DONOHO, D. L. and JOHNSTONE, I. M. (1994). Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81, 425–455. MR1311089
- [6] DURRETT, R. (1996). *Probability: Theory and Examples*. Second edition. Duxbury Press, Belmont, CA. MR1609153

- [7] FAN, J. and LIN, S. K. (1998). Test of significance when data are curves. *Journal of American Statistical Association*, **93**, 1007–1021. MR1649196
- [8] FUH, C.D. and MEI, Y. (2015). Quickest change detection and Kullback-Leibler divergence for two-state hidden Markov models. *IEEE Trans. Signal Processing*, **63**, 4866–4878.
- [9] GLAZ, J., NAUS, J. and WALLENSTEIN, S. (2001). *Scan Statistics*. Springer-Verlag, New York. MR1869112
- [10] GORDON, L. and POLLAK, M. (1994). An efficient sequential nonparametric scheme for detecting a change of distribution. *Ann. Statist.* **22** 763–804. MR1292540
- [11] KIEFER, J. and SACKS, J. (1963). Asymptotically optimum sequential inference and design. *Ann. Math. Statist.* **34** 705–750. MR0150907
- [12] KULLDORFF, M. (2001). Prospective Time-Periodic Geographic Disease Surveillance Using a Scan Statistic, *J. R. Stat. Soc. Ser. A* **164** 61–72. MR1819022
- [13] LAI, T. L. (1995). Sequential change-point detection in quality control and dynamical systems (with discussion). *J. R. Stat. Soc. Ser. B Stat. Methodol.* **57** 613–658. MR1354072
- [14] LAI, T. L. (2001). Sequential analysis: some classical problems and new challenges. *Statist. Sinica* **11** 303–408. MR1844531
- [15] LÉVY-LEDUC, C. and ROUEFF, F. (2009). Detection and localization of change-points in high-dimensional network traffic data. *Ann. Appl. Stat.* **3** 637–662. MR2750676
- [16] LIU, K., MEI, Y., and SHI, J. (2015). An adaptive sampling strategy for online high-dimensional process monitoring. *Technometrics* **57** 305–319. MR3384946
- [17] LORDEN, G. (1971). Procedures for reacting to a change in distribution. *Ann. Math. Statist.* **42** 1897–1908. MR0309251
- [18] LORDEN, G. and POLLAK, M. (2008). Sequential change-point detection procedures that are nearly optimal and computationally simple. *Sequential Analysis* **27** 476–512. MR2460209
- [19] MEI, Y. (2005). Information bounds and quickest change detection in decentralized decision systems. *IEEE Trans. Inform. Theory* **51** 2669–2681. MR2246385
- [20] MEI, Y. (2010). Efficient scalable schemes for monitoring a large number of data streams. *Biometrika* **97.2** 419–433. MR2650748
- [21] MONTGOMERY, D. C. (1991). *Introduction to Statistical Quality Control* (2nd edition). Wiley, New York.

- [22] MOUSTAKIDES, G. V. (1986). Optimal stopping times for detecting changes in distributions. *Ann. Statist.* **14** 1379–1387. MR0868306
- [23] NEYMAN, J. (1937). Smooth test for goodness-of-fit. *Skand. Aktuarietidskr.* **20** 149–199.
- [24] PAGE, E. S. (1954). Continuous inspection schemes. *Biometrika* **41** 100–115. MR0088850
- [25] POLLAK, M. (1985). Optimal detection of a change in distribution. *Ann. Statist.* **13** 206–227. MR0773162
- [26] POLLAK, M. (1987). Average run lengths of an optimal method of detecting a change in distribution. *Ann. Statist.* **15** 749–779. MR0888438
- [27] POOR, H. V. and HADJILIADIS, O. (2009). *Quickest Detection*. Cambridge Univ. Press, New York, 2009. MR2482527
- [28] RAGO, C., WILLETT, P., and BAR-SHALOM, Y. (1996). Censoring sensors: A low-communication-rate scheme for distributed detection. *IEEE Trans. Aerosp. Electron. Syst.*, *32*, 554–568.
- [29] ROBERTS, S. W. (1966). A comparison of some control chart procedures. *Technometrics* **8** 411–430. MR0196887
- [30] SHEWHART, W. A. (1931). *Economic Control of Quality of Manufactured Product*. D Van Norstrand, New York. Preprinted by ASQC Quality Press, Wisconsin, 1980.
- [31] SHIRYAEV, A. N. (1963). On optimum methods in quickest detection problems. *Theory Probab. Appl.* **8** 22–46.
- [32] SIEGMUND, D. (1985): *Sequential Analysis: Tests and Confidence Intervals*. Springer, New York. MR0799155
- [33] TARTAKOVSKY, A., NIKIFOROV, I., and BASSEVILLE, M. (2015). *Sequential Analysis: Hypothesis Testing and Changepoint Detection*. Monographs on Statistics and Applied Probability, 136. CRC Press, Boca Raton, FL. MR3241619
- [34] TARTAKOVSKY, A. G., ROZOVSKIIA, B. L., BLAZEKA, R. B. and KIM, H. (2006). Detection of intrusions in information systems by sequential change-point methods (with discussions). *Statistical Methodology* **3** 252–340. MR2240956
- [35] TARTAKOVSKY, A. G. and VEERAVALLI, V. V. (2004). Change-point Detection in Multichannel and Distributed Systems. *Applied Sequential Methodologies*, 339–370, Statist. Textbooks Monogr., 173, Dekker, New York. MR2159163

- [36] TAY, W. P., TSITSIKLIS, J. N. and WIN, M. Z. (2007). Asymptotic performance of a censoring sensor network. *IEEE Trans. Inform. Theory* **53** 4191–4209. MR2446562
- [37] XIE, Y., HUANG, J., and WILLETT, R. (2013). Change-point detection for high-dimensional time series with missing data. *IEEE Journal of Selected Topics in Signal Processing*, **7**, 12–27.
- [38] XIE, Y. and SIEGMUND, D. (2013). Sequential multi-sensor change-point detection. *Ann. Stat.*, **41** 670–692. MR3099117
- [39] VEERAVALLI, V. V. (2001). Decentralized quickest change detection. *IEEE Trans. Inform. Theory* **47** 1657–1665. MR1830119